
**Information technology — Database
languages — SQL —**

**Part 2:
Foundation (SQL/Foundation)**

*Technologies de l'information — Langages de base de données —
SQL —*

Partie 2: Fondations (SQL/Foundation)

PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2008

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents	Page
Foreword.....	xix
Introduction.....	xx
1 Scope.....	1
2 Normative references.....	3
2.1 ISO and IEC standards.....	3
2.2 Other international standards.....	4
3 Definitions, notations, and conventions.....	5
3.1 Definitions.....	5
3.1.1 Definitions taken from [ISO10646].....	5
3.1.2 Definitions taken from [ISO14651].....	5
3.1.3 Definitions taken from [Unicode].....	5
3.1.4 Definitions taken from [ISO8601].....	6
3.1.5 Definitions taken from [ISO9075-1].....	6
3.1.6 Definitions provided in Part 2.....	7
3.2 Notation.....	11
3.3 Conventions.....	11
3.3.1 Use of terms.....	11
3.3.1.1 Other terms.....	11
4 Concepts.....	13
4.1 Data types.....	13
4.1.1 General introduction to data types.....	13
4.1.2 Naming of predefined types.....	14
4.1.3 Host language data types.....	15
4.1.4 Data type terminology.....	15
4.1.5 Properties of distinct.....	17
4.2 Character strings.....	18
4.2.1 Introduction to character strings.....	18
4.2.2 Comparison of character strings.....	19
4.2.3 Operations involving character strings.....	19
4.2.3.1 Regular expression syntaxes.....	19
4.2.3.2 Operators that operate on character strings and return character strings.....	20
4.2.3.3 Other operators involving character strings.....	22
4.2.3.4 Operations involving large object character strings.....	23
4.2.4 Character repertoires.....	24
4.2.5 Character encoding forms.....	24

4.2.6	Collations.	25
4.2.7	Character sets.	26
4.2.8	Universal character sets.	28
4.3	Binary strings.	28
4.3.1	Introduction to binary strings.	28
4.3.2	Binary string comparison.	28
4.3.3	Operations involving binary strings.	29
4.3.3.1	Operators that operate on binary strings and return binary strings.	29
4.3.3.2	Other operators involving binary strings.	29
4.4	Numbers.	30
4.4.1	Introduction to numbers.	30
4.4.2	Characteristics of numbers.	30
4.4.3	Operations involving numbers.	31
4.5	Boolean types.	33
4.5.1	Introduction to Boolean types.	33
4.5.2	Comparison and assignment of booleans.	33
4.5.3	Operations involving booleans.	33
4.5.3.1	Operations on booleans that return booleans.	33
4.5.3.2	Other operators involving booleans.	33
4.6	Datetimes and intervals.	34
4.6.1	Introduction to datetimes and intervals.	34
4.6.2	Datetimes.	34
4.6.3	Intervals.	37
4.6.4	Operations involving datetimes and intervals.	39
4.7	User-defined types.	40
4.7.1	Introduction to user-defined types.	40
4.7.2	User-defined type descriptor.	41
4.7.3	Observer functions and mutator functions.	42
4.7.4	Constructors.	43
4.7.5	Subtypes and supertypes.	43
4.7.6	User-defined type comparison and assignment.	44
4.7.7	Transforms for user-defined types.	45
4.8	Row types.	46
4.9	Reference types.	46
4.9.1	Introduction to reference types.	46
4.9.2	Operations involving references.	47
4.10	Collection types.	47
4.10.1	Introduction to collection types.	47
4.10.2	Arrays.	48
4.10.3	Multisets.	48
4.10.4	Collection comparison and assignment.	49
4.10.5	Operations involving arrays.	49
4.10.5.1	Operators that operate on array values and return array elements.	49
4.10.5.2	Operators that operate on array values and return array values.	49

4.10.5.3	Operators that operate on array values and return numbers.	49
4.10.6	Operations involving multisets.	50
4.10.6.1	Operators that operate on multisets and return multiset elements.	50
4.10.6.2	Operators that operate on multisets and return multisets.	50
4.10.6.3	Operators that operate on multiset values and return numbers.	50
4.11	Data conversions.	50
4.12	Domains.	51
4.13	Columns, fields, and attributes.	52
4.14	Tables.	53
4.14.1	Introduction to tables.	53
4.14.2	Types of tables.	54
4.14.3	Table descriptors.	56
4.14.4	Relationships between tables.	57
4.14.5	Referenceable tables, subtables, and supertables.	58
4.14.6	Operations involving tables.	59
4.14.7	Identity columns.	60
4.14.8	Base columns and generated columns.	61
4.14.9	Windowed tables.	61
4.15	Data analysis operations (involving tables).	62
4.15.1	Introduction to data analysis operations.	62
4.15.2	Group functions.	63
4.15.3	Window functions.	63
4.15.4	Aggregate functions.	64
4.16	Determinism.	66
4.17	Integrity constraints.	67
4.17.1	Overview of integrity constraints.	67
4.17.2	Checking of constraints.	67
4.17.3	Table constraints.	68
4.17.4	Domain constraints.	70
4.17.5	Assertions.	71
4.18	Functional dependencies.	71
4.18.1	Overview of functional dependency rules and notations.	71
4.18.2	General rules and definitions.	72
4.18.3	Known functional dependencies in a base table.	73
4.18.4	Known functional dependencies in a transition table.	74
4.18.5	Known functional dependencies in <table value constructor>.	74
4.18.6	Known functional dependencies in a <joined table>.	74
4.18.7	Known functional dependencies in a <table primary>.	76
4.18.8	Known functional dependencies in a <table factor>.	76
4.18.9	Known functional dependencies in a <table reference>.	77
4.18.10	Known functional dependencies in the result of a <from clause>.	77
4.18.11	Known functional dependencies in the result of a <where clause>.	77
4.18.12	Known functional dependencies in the result of a <group by clause>.	78
4.18.13	Known functional dependencies in the result of a <having clause>.	78

4.18.14	Known functional dependencies in a <query specification>.....	78
4.18.15	Known functional dependencies in a <query expression>.....	79
4.19	Candidate keys.....	80
4.20	SQL-schemas.....	81
4.21	Sequence generators.....	82
4.21.1	General description of sequence generators.....	82
4.21.2	Operations involving sequence generators.....	83
4.22	SQL-client modules.....	83
4.23	Embedded syntax.....	84
4.24	Dynamic SQL concepts.....	85
4.24.1	Overview of dynamic SQL.....	85
4.24.2	Dynamic SQL statements and descriptor areas.....	86
4.25	Direct invocation of SQL.....	87
4.26	Externally-invoked procedures.....	87
4.27	SQL-invoked routines.....	88
4.27.1	Overview of SQL-invoked routines.....	88
4.27.2	Characteristics of SQL-invoked routines.....	89
4.27.3	Execution of SQL-invoked routines.....	91
4.27.4	Routine descriptors.....	91
4.27.5	Result sets returned by SQL-invoked procedures.....	93
4.28	SQL-paths.....	94
4.29	Host parameters.....	95
4.29.1	Overview of host parameters.....	95
4.29.2	Status parameters.....	95
4.29.3	Data parameters.....	96
4.29.4	Indicator parameters.....	96
4.29.5	Locators.....	96
4.30	Diagnostics area.....	97
4.31	Host languages.....	98
4.32	Cursors.....	99
4.32.1	General description of cursors.....	99
4.32.2	Operations on and using cursors.....	102
4.33	SQL-statements.....	104
4.33.1	Classes of SQL-statements.....	104
4.33.2	SQL-statements classified by function.....	105
4.33.2.1	SQL-schema statements.....	105
4.33.2.2	SQL-data statements.....	106
4.33.2.3	SQL-data change statements.....	107
4.33.2.4	SQL-transaction statements.....	107
4.33.2.5	SQL-connection statements.....	108
4.33.2.6	SQL-control statements.....	108
4.33.2.7	SQL-session statements.....	108
4.33.2.8	SQL-diagnostics statements.....	109
4.33.2.9	SQL-dynamic statements.....	109

4.33.2.10	SQL embedded exception declaration.	109
4.33.3	SQL-statements and SQL-data access indication.	110
4.33.4	SQL-statements and transaction states.	110
4.33.5	SQL-statement atomicity and statement execution contexts.	112
4.33.6	Embeddable SQL-statements.	113
4.33.7	Preparable and immediately executable SQL-statements.	115
4.33.8	Directly executable SQL-statements.	116
4.34	Basic security model.	118
4.34.1	Authorization identifiers.	118
4.34.1.1	SQL-session authorization identifiers.	118
4.34.1.2	SQL-client module authorization identifiers.	119
4.34.1.3	SQL-schema authorization identifiers.	119
4.34.2	Privileges.	119
4.34.3	Roles.	121
4.34.4	Security model definitions.	122
4.35	SQL-transactions.	122
4.35.1	General description of SQL-transactions.	122
4.35.2	Savepoints.	123
4.35.3	Properties of SQL-transactions.	124
4.35.4	Isolation levels of SQL-transactions.	124
4.35.5	Implicit rollbacks.	126
4.35.6	Effects of SQL-statements in an SQL-transaction.	126
4.35.7	Encompassing transactions.	126
4.36	SQL-connections.	127
4.37	SQL-sessions.	128
4.37.1	General description of SQL-sessions.	128
4.37.2	SQL-session identification.	129
4.37.3	SQL-session properties.	130
4.37.4	Execution contexts.	132
4.37.5	Routine execution context.	132
4.38	Triggers.	133
4.38.1	General description of triggers.	133
4.38.2	Trigger execution.	135
4.39	Client-server operation.	137
5	Lexical elements.	139
5.1	<SQL terminal character>.	139
5.2	<token> and <separator>.	143
5.3	<literal>.	152
5.4	Names and identifiers.	161
6	Scalar expressions.	173
6.1	<data type>.	173
6.2	<field definition>.	186
6.3	<value expression primary>.	188

6.4	<value specification> and <target specification>.....	190
6.5	<contextually typed value specification>.....	196
6.6	<identifier chain>.....	198
6.7	<column reference>.....	202
6.8	<SQL parameter reference>.....	204
6.9	<set function specification>.....	205
6.10	<window function>.....	207
6.11	<case expression>.....	211
6.12	<cast specification>.....	215
6.13	<next value expression>.....	230
6.14	<field reference>.....	232
6.15	<subtype treatment>.....	233
6.16	<method invocation>.....	235
6.17	<static method invocation>.....	237
6.18	<new specification>.....	239
6.19	<attribute or method reference>.....	241
6.20	<dereference operation>.....	243
6.21	<method reference>.....	244
6.22	<reference resolution>.....	246
6.23	<array element reference>.....	248
6.24	<multiset element reference>.....	249
6.25	<value expression>.....	250
6.26	<numeric value expression>.....	254
6.27	<numeric value function>.....	256
6.28	<string value expression>.....	268
6.29	<string value function>.....	273
6.30	<datetime value expression>.....	288
6.31	<datetime value function>.....	291
6.32	<interval value expression>.....	293
6.33	<interval value function>.....	298
6.34	<boolean value expression>.....	299
6.35	<array value expression>.....	304
6.36	<array value constructor>.....	306
6.37	<multiset value expression>.....	308
6.38	<multiset value function>.....	311
6.39	<multiset value constructor>.....	312
7	Query expressions.....	315
7.1	<row value constructor>.....	315
7.2	<row value expression>.....	318
7.3	<table value constructor>.....	320
7.4	<table expression>.....	322
7.5	<from clause>.....	323
7.6	<table reference>.....	325

7.7	<joined table>.....	335
7.8	<where clause>.....	344
7.9	<group by clause>.....	345
7.10	<having clause>.....	354
7.11	<window clause>.....	356
7.12	<query specification>.....	366
7.13	<query expression>.....	376
7.14	<search or cycle clause>.....	393
7.15	<subquery>.....	398
8	Predicates.....	401
8.1	<predicate>.....	401
8.2	<comparison predicate>.....	403
8.3	<between predicate>.....	411
8.4	<in predicate>.....	412
8.5	<like predicate>.....	414
8.6	<similar predicate>.....	420
8.7	<regex like predicate>.....	426
8.8	<null predicate>.....	428
8.9	<quantified comparison predicate>.....	430
8.10	<exists predicate>.....	432
8.11	<unique predicate>.....	433
8.12	<normalized predicate>.....	434
8.13	<match predicate>.....	436
8.14	<overlaps predicate>.....	439
8.15	<distinct predicate>.....	441
8.16	<member predicate>.....	444
8.17	<submultiset predicate>.....	446
8.18	<set predicate>.....	448
8.19	<type predicate>.....	449
8.20	<search condition>.....	451
9	Additional common rules.....	453
9.1	Retrieval assignment.....	453
9.2	Store assignment.....	458
9.3	Result of data type combinations.....	463
9.4	Subject routine determination.....	466
9.5	Type precedence list determination.....	468
9.6	Host parameter mode determination.....	472
9.7	Type name determination.....	474
9.8	Determination of identical values.....	476
9.9	Equality operations.....	478
9.10	Grouping operations.....	481
9.11	Multiset element grouping operations.....	483
9.12	Ordering operations.....	485

9.13	Collation determination.	487
9.14	Execution of array-returning functions.	489
9.15	Execution of multiset-returning functions.	492
9.16	XQuery regular expression matching.	493
9.17	XQuery regular expression replacement.	496
9.18	Data type identity.	498
9.19	Determination of a from-sql function.	500
9.20	Determination of a from-sql function for an overriding method.	501
9.21	Determination of a to-sql function.	502
9.22	Determination of a to-sql function for an overriding method.	503
9.23	Generation of the next value of a sequence generator.	504
9.24	Creation of a sequence generator.	505
9.25	Altering a sequence generator.	507
10	Additional common elements.	509
10.1	<interval qualifier>.	509
10.2	<language clause>.	513
10.3	<path specification>.	515
10.4	<routine invocation>.	516
10.5	<character set specification>.	539
10.6	<specific routine designator>.	541
10.7	<collate clause>.	544
10.8	<constraint name definition> and <constraint characteristics>.	545
10.9	<aggregate function>.	547
10.10	<sort specification list>.	559
10.11	Determination of view privileges.	561
10.12	Determination of view component privileges.	563
11	Schema definition and manipulation.	567
11.1	<schema definition>.	567
11.2	<drop schema statement>.	570
11.3	<table definition>.	573
11.4	<column definition>.	584
11.5	<default clause>.	589
11.6	<table constraint definition>.	593
11.7	<unique constraint definition>.	595
11.8	<referential constraint definition>.	597
11.9	<check constraint definition>.	601
11.10	<alter table statement>.	603
11.11	<add column definition>.	604
11.12	<alter column definition>.	606
11.13	<set column default clause>.	607
11.14	<drop column default clause>.	608
11.15	<add column scope clause>.	609
11.16	<drop column scope clause>.	610

11.17	<alter column data type clause>.....	612
11.18	<alter identity column specification>.....	615
11.19	<drop column definition>.....	616
11.20	<add table constraint definition>.....	618
11.21	<drop table constraint definition>.....	619
11.22	<drop table statement>.....	622
11.23	<view definition>.....	625
11.24	<drop view statement>.....	635
11.25	<domain definition>.....	638
11.26	<alter domain statement>.....	641
11.27	<set domain default clause>.....	642
11.28	<drop domain default clause>.....	643
11.29	<add domain constraint definition>.....	644
11.30	<drop domain constraint definition>.....	645
11.31	<drop domain statement>.....	646
11.32	<character set definition>.....	648
11.33	<drop character set statement>.....	650
11.34	<collation definition>.....	652
11.35	<drop collation statement>.....	654
11.36	<transliteration definition>.....	656
11.37	<drop transliteration statement>.....	659
11.38	<assertion definition>.....	661
11.39	<drop assertion statement>.....	663
11.40	<trigger definition>.....	666
11.41	<drop trigger statement>.....	671
11.42	<user-defined type definition>.....	674
11.43	<attribute definition>.....	690
11.44	<alter type statement>.....	692
11.45	<add attribute definition>.....	693
11.46	<drop attribute definition>.....	695
11.47	<add original method specification>.....	697
11.48	<add overriding method specification>.....	703
11.49	<drop method specification>.....	708
11.50	<drop data type statement>.....	712
11.51	<SQL-invoked routine>.....	715
11.52	<alter routine statement>.....	740
11.53	<drop routine statement>.....	743
11.54	<user-defined cast definition>.....	745
11.55	<drop user-defined cast statement>.....	747
11.56	<user-defined ordering definition>.....	749
11.57	<drop user-defined ordering statement>.....	753
11.58	<transform definition>.....	755
11.59	<alter transform statement>.....	758
11.60	<add transform element list>.....	759

11.61	<drop transform element list>.....	761
11.62	<drop transform statement>.....	763
11.63	<sequence generator definition>.....	766
11.64	<alter sequence generator statement>.....	768
11.65	<drop sequence generator statement>.....	769
12	Access control.....	771
12.1	<grant statement>.....	771
12.2	<grant privilege statement>.....	776
12.3	<privileges>.....	779
12.4	<role definition>.....	782
12.5	<grant role statement>.....	783
12.6	<drop role statement>.....	785
12.7	<revoke statement>.....	786
12.8	Grantor determination.....	804
13	SQL-client modules.....	805
13.1	<SQL-client module definition>.....	805
13.2	<module name clause>.....	810
13.3	<externally-invoked procedure>.....	811
13.4	Calls to an <externally-invoked procedure>.....	814
13.5	<SQL procedure statement>.....	831
13.6	Data type correspondences.....	838
14	Data manipulation.....	849
14.1	<declare cursor>.....	849
14.2	<cursor properties>.....	851
14.3	<cursor specification>.....	853
14.4	<open statement>.....	856
14.5	<fetch statement>.....	857
14.6	<close statement>.....	861
14.7	<select statement: single row>.....	862
14.8	<delete statement: positioned>.....	866
14.9	<delete statement: searched>.....	868
14.10	<truncate table statement>.....	871
14.11	<insert statement>.....	873
14.12	<merge statement>.....	878
14.13	<update statement: positioned>.....	886
14.14	<update statement: searched>.....	888
14.15	<set clause list>.....	892
14.16	<temporary table declaration>.....	897
14.17	<free locator statement>.....	899
14.18	<hold locator statement>.....	900
15	Additional data manipulation rules.....	901
15.1	Effect of opening a cursor.....	901
15.2	Effect of receiving a result set.....	904

15.3	Determination of the current row of a cursor.	905
15.4	Effect of closing a cursor.	907
15.5	Effect of a positioned delete.	908
15.6	Effect of a positioned update.	910
15.7	Effect of deleting rows from base tables.	913
15.8	Effect of deleting some rows from a derived table.	915
15.9	Effect of deleting some rows from a viewed table.	917
15.10	Effect of inserting tables into base tables.	919
15.11	Effect of inserting a table into a derived table.	921
15.12	Effect of inserting a table into a viewed table.	923
15.13	Effect of replacing rows in base tables.	925
15.14	Effect of replacing some rows in a derived table.	927
15.15	Effect of replacing some rows in a viewed table.	930
15.16	Execution of BEFORE triggers.	932
15.17	Execution of referential actions.	933
15.18	Execution of AFTER triggers.	939
15.19	Execution of triggers.	940
16	Control statements.	943
16.1	<call statement>.	943
16.2	<return statement>.	944
17	Transaction management.	945
17.1	<start transaction statement>.	945
17.2	<set transaction statement>.	947
17.3	<transaction characteristics>.	949
17.4	<set constraints mode statement>.	951
17.5	<savepoint statement>.	953
17.6	<release savepoint statement>.	954
17.7	<commit statement>.	955
17.8	<rollback statement>.	957
18	Connection management.	961
18.1	<connect statement>.	961
18.2	<set connection statement>.	964
18.3	<disconnect statement>.	966
19	Session management.	969
19.1	<set session characteristics statement>.	969
19.2	<set session user identifier statement>.	971
19.3	<set role statement>.	972
19.4	<set local time zone statement>.	973
19.5	<set catalog statement>.	974
19.6	<set schema statement>.	975
19.7	<set names statement>.	977
19.8	<set path statement>.	978
19.9	<set transform group statement>.	979

19.10	<set session collation statement>.....	980
20	Dynamic SQL.....	983
20.1	Description of SQL descriptor areas.....	983
20.2	<allocate descriptor statement>.....	993
20.3	<deallocate descriptor statement>.....	995
20.4	<get descriptor statement>.....	996
20.5	<set descriptor statement>.....	999
20.6	<prepare statement>.....	1003
20.7	<cursor attributes>.....	1014
20.8	<deallocate prepared statement>.....	1015
20.9	<describe statement>.....	1017
20.10	<input using clause>.....	1023
20.11	<output using clause>.....	1027
20.12	<execute statement>.....	1032
20.13	<execute immediate statement>.....	1034
20.14	<dynamic declare cursor>.....	1035
20.15	<allocate cursor statement>.....	1036
20.16	<dynamic open statement>.....	1039
20.17	<dynamic fetch statement>.....	1041
20.18	<dynamic single row select statement>.....	1042
20.19	<dynamic close statement>.....	1043
20.20	<dynamic delete statement: positioned>.....	1044
20.21	<dynamic update statement: positioned>.....	1046
20.22	<preparable dynamic delete statement: positioned>.....	1048
20.23	<preparable dynamic cursor name>.....	1050
20.24	<preparable dynamic update statement: positioned>.....	1052
21	Embedded SQL.....	1055
21.1	<embedded SQL host program>.....	1055
21.2	<embedded exception declaration>.....	1067
21.3	<embedded SQL Ada program>.....	1071
21.4	<embedded SQL C program>.....	1078
21.5	<embedded SQL COBOL program>.....	1087
21.6	<embedded SQL Fortran program>.....	1094
21.7	<embedded SQL MUMPS program>.....	1100
21.8	<embedded SQL Pascal program>.....	1105
21.9	<embedded SQL PL/I program>.....	1111
22	Direct invocation of SQL.....	1119
22.1	<direct SQL statement>.....	1119
22.2	<direct select statement: multiple rows>.....	1123
23	Diagnostics management.....	1125
23.1	<get diagnostics statement>.....	1125
23.2	Pushing and popping the diagnostics area stack.....	1142

24	Status codes.....	1143
24.1	SQLSTATE.....	1143
24.2	Remote Database Access SQLSTATE Subclasses.....	1152
25	Conformance.....	1153
25.1	Claims of conformance to SQL/Foundation.....	1153
25.2	Additional conformance requirements for SQL/Foundation.....	1154
25.3	Implied feature relationships of SQL/Foundation.....	1154
Annex A	SQL Conformance Summary (informative).....	1159
Annex B	Implementation-defined elements (informative).....	1225
Annex C	Implementation-dependent elements (informative).....	1245
Annex D	Deprecated features (informative).....	1253
Annex E	Incompatibilities with ISO/IEC 9075:2003 (informative).....	1255
Annex F	SQL feature taxonomy (informative).....	1257
Annex G	Defect Reports not addressed in this edition of ISO/IEC 9075 (informative).....	1287
	Bibliography.....	1289
	Index.....	1291

Tables

Table	Page
1 Overview of character sets.	27
2 Fields in datetime values.	35
3 Datetime data type conversions.	36
4 Fields in year-month INTERVAL values.	37
5 Fields in day-time INTERVAL values.	38
6 Valid values for fields in INTERVAL values.	38
7 Valid operators involving datetimes and intervals.	39
8 SQL-transaction isolation levels and the three phenomena.	125
9 Valid values for datetime fields.	181
10 Valid absolute values for interval fields.	182
11 Truth table for the AND boolean operator.	302
12 Truth table for the OR boolean operator.	302
13 Truth table for the IS boolean operator.	302
14 <null predicate> semantics.	429
15 Standard programming languages.	513
16 Data type correspondences for Ada.	838
17 Data type correspondences for C.	839
18 Data type correspondences for COBOL.	841
19 Data type correspondences for Fortran.	842
20 Data type correspondences for M.	844
21 Data type correspondences for Pascal.	845
22 Data type correspondences for PL/I.	846
23 Data types of <key word>s used in the header of SQL descriptor areas.	987
24 Data types of <key word>s used in SQL item descriptor areas.	987
25 Codes used for SQL data types in Dynamic SQL.	989
26 Codes associated with datetime data types in Dynamic SQL.	990
27 Codes used for <interval qualifier>s in Dynamic SQL.	991
28 Codes used for input/output SQL parameter modes in Dynamic SQL.	992
29 Codes associated with user-defined types in Dynamic SQL.	992
30 <statement information item name>s for use with <get diagnostics statement>.	1127
31 <condition information item name>s for use with <get diagnostics statement>.	1127
32 SQL-statement codes.	1130
33 SQLSTATE class and subclass values.	1144
34 SQLSTATE class codes for RDA.	1152
35 Implied feature relationships of SQL/Foundation.	1154
36 Feature taxonomy and definition for mandatory features.	1257
37 Feature taxonomy for optional features.	1274

Figures

Figure	Page
1 Operation of <regular expression substring function>.....	21
2 Illustration of WIDTH_BUCKET Semantics.....	32

(Blank page)

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75% of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

International Standard ISO/IEC 9075-2 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 32, *Data management and interchange*.

This third edition of ISO/IEC 9075-2 cancels and replaces the second edition (ISO/IEC 9075-2:2003), which has been technically revised. It also incorporates Technical Corrigendum ISO/IEC 9075-2:2003/Cor.2:2007.

ISO/IEC 9075 consists of the following parts, under the general title *Information technology — Database languages — SQL*:

- Part 1: Framework (SQL/Framework)
- Part 2: Foundation (SQL/Foundation)
- Part 3: Call-Level Interface (SQL/CLI)
- Part 4: Persistent Stored Modules (SQL/PSM)
- Part 9: Management of External Data (SQL/MED)
- Part 10: Object Language Bindings (SQL/OLB)
- Part 11: Information and Definition Schema (SQL/Schemata)
- Part 13: SQL Routines and Types Using the Java™ Programming Language (SQL/JRT)
- Part 14: XML-Related Specifications (SQL/XML)

Introduction

The organization of this part of ISO/IEC 9075 is as follows:

- 1) [Clause 1, “Scope”](#), specifies the scope of this part of ISO/IEC 9075.
- 2) [Clause 2, “Normative references”](#), identifies additional standards that, through reference in this part of ISO/IEC 9075, constitute provisions of this part of ISO/IEC 9075.
- 3) [Clause 3, “Definitions, notations, and conventions”](#), defines the notations and conventions used in this part of ISO/IEC 9075.
- 4) [Clause 4, “Concepts”](#), presents concepts used in the definition of SQL.
- 5) [Clause 5, “Lexical elements”](#), defines the lexical elements of the language.
- 6) [Clause 6, “Scalar expressions”](#), defines the elements of the language that produce scalar values.
- 7) [Clause 7, “Query expressions”](#), defines the elements of the language that produce rows and tables of data.
- 8) [Clause 8, “Predicates”](#), defines the predicates of the language.
- 9) [Clause 9, “Additional common rules”](#), specifies the rules for assignments that retrieve data from or store data into SQL-data, and formation rules for set operations.
- 10) [Clause 10, “Additional common elements”](#), defines additional language elements that are used in various parts of the language.
- 11) [Clause 11, “Schema definition and manipulation”](#), defines facilities for creating and managing a schema.
- 12) [Clause 12, “Access control”](#), defines facilities for controlling access to SQL-data.
- 13) [Clause 13, “SQL-client modules”](#), defines SQL-client modules and externally-invoked procedures.
- 14) [Clause 14, “Data manipulation”](#), defines the data manipulation statements.
- 15) [Clause 15, “Additional data manipulation rules”](#), defines additional rules for data manipulation.
- 16) [Clause 16, “Control statements”](#), defines the SQL-control statements.
- 17) [Clause 17, “Transaction management”](#), defines the SQL-transaction management statements.
- 18) [Clause 18, “Connection management”](#) defines the SQL-connection management statements.
- 19) [Clause 19, “Session management”](#), defines the SQL-session management statements.
- 20) [Clause 20, “Dynamic SQL”](#), defines the SQL dynamic statements.
- 21) [Clause 21, “Embedded SQL”](#), defines the host language embeddings.
- 22) [Clause 22, “Direct invocation of SQL”](#), defines direct invocation of SQL language.
- 23) [Clause 23, “Diagnostics management”](#), defines the diagnostics management facilities.
- 24) [Clause 24, “Status codes”](#), defines values that identify the status of the execution of SQL-statements and the mechanisms by which those values are returned.

- 25) [Clause 25, “Conformance”](#), defines the criteria for conformance to this part of ISO/IEC 9075.
- 26) [Annex A, “SQL Conformance Summary”](#), is an informative Annex. It summarizes the conformance requirements of the SQL language.
- 27) [Annex B, “Implementation-defined elements”](#), is an informative Annex. It lists those features for which the body of this part of ISO/IEC 9075 states that the syntax, the meaning, the returned results, the effect on SQL-data and/or schemas, or any other behavior is partly or wholly implementation-defined.
- 28) [Annex C, “Implementation-dependent elements”](#), is an informative Annex. It lists those features for which the body of this part of ISO/IEC 9075 states that the syntax, the meaning, the returned results, the effect on SQL-data and/or schemas, or any other behavior is partly or wholly implementation-dependent.
- 29) [Annex D, “Deprecated features”](#), is an informative Annex. It lists features that the responsible Technical Committee intend will not appear in a future revised version of this part of ISO/IEC 9075.
- 30) [Annex E, “Incompatibilities with ISO/IEC 9075:2003”](#), is an informative Annex. It lists incompatibilities with the previous version of this part of ISO/IEC 9075.
- 31) [Annex F, “SQL feature taxonomy”](#), is an informative Annex. It identifies features of the SQL language specified in this part of ISO/IEC 9075 by an identifier and a short descriptive name. This taxonomy is used to specify conformance and may be used to develop other profiles involving the SQL language.
- 32) [Annex G, “Defect Reports not addressed in this edition of ISO/IEC 9075”](#), is an informative Annex. It describes the Defect Reports that were known at the time of publication of this part of this International Standard. Each of these problems is a problem carried forward from the previous edition of ISO/IEC 9075. No new problems have been created in the drafting of this edition of this International Standard.

In the text of this part of ISO/IEC 9075, Clauses begin a new odd-numbered page, and in [Clause 5, “Lexical elements”](#), through [Clause 24, “Status codes”](#), Subclauses begin a new page. Any resulting blank space is not significant.

Information technology — Database languages — SQL —

Part 2:

Foundation (SQL/Foundation)

1 Scope

This part of ISO/IEC 9075 defines the data structures and basic operations on SQL-data. It provides functional capabilities for creating, accessing, maintaining, controlling, and protecting SQL-data.

This part of ISO/IEC 9075 specifies the syntax and semantics of a database language:

- For specifying and modifying the structure and the integrity constraints of SQL-data.
- For declaring and invoking operations on SQL-data and cursors.
- For declaring database language procedures.
- For embedding SQL-statements in a compilation unit that is otherwise written in a particular programming language (host language).
- For deriving an equivalent compilation unit in the host language. In that equivalent compilation unit, each embedded SQL-statement has been replaced by one or more statements in the host language, some of which invoke an SQL externally-invoked procedure that, when executed, has an effect equivalent to executing the SQL-statement.
- For direct invocation of SQL-statements.
- To support dynamic preparation and execution of SQL-statements.

This part of ISO/IEC 9075 provides a vehicle for portability of data definitions and compilation units between SQL-implementations.

This part of ISO/IEC 9075 provides a vehicle for interconnection of SQL-implementations.

Implementations of this part of ISO/IEC 9075 may exist in environments that also support application programming languages, end-user query languages, report generator systems, data dictionary systems, program library systems, and distributed communication systems, as well as various tools for database design, data administration, and performance optimization.

(Blank page)

2 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

2.1 ISO and IEC standards

[ISO646] ISO/IEC 646:1991, *Information technology — ISO 7-bit coded character set for information interchange*.

[ISO1539] ISO/IEC 1539-1:2004, *Information technology — Programming languages — Fortran — Part 1: Base language*.

[ISO1989] ISO/IEC 1989:2002, *Information technology — Programming languages — COBOL*.

[ISO6160] ISO 6160:1979, *Programming languages — PL/I* (Endorsement of ANSI X3.53-1976).

[ISO6429] ISO/IEC 6429:1992, *Information technology — Control functions for coded character sets*

[ISO7185] ISO 7185:1990, *Information technology — Programming languages — Pascal*.

[ISO8601] ISO 8601:2004, *Data elements and interchange formats — Information interchange — Representation of dates and times*.

[ISO8649] ISO/IEC 8649:1996, *Information technology — Open Systems Interconnection — Service Definition for the Association Control Service Element*.

ISO/IEC 8649:1996/Amd.1:1997, *Support of authentication mechanisms for the connectionless mode*.

ISO/IEC 8649:1996/Amd.2:1998, *Fast-associate mechanism*.

[ISO8652] ISO/IEC 8652:1995, *Information technology — Programming languages — Ada*.

ISO/IEC 8652:1995/Cor.1:2001.

[ISO8859-1] ISO/IEC 8859-1:1998, *Information technology — 8-bit single-byte coded graphic character sets — Part 1: Latin alphabet No. 1*

[ISO9075-1] ISO/IEC 9075-1:2008, *Information technology — Database languages — SQL — Part 1: Framework (SQL/Framework)*.

[ISO9075-11] ISO/IEC 9075-11:2008, *Information technology — Database languages — SQL — Part 11: Information and Definition Schemas (SQL/Schemata)*.

[ISO9579] ISO/IEC 9579:2000, *Information technology — Remote database access for SQL with security enhancement*.

[ISO9899] ISO/IEC 9899:1999, *Programming languages — C*.

ISO/IEC 9899:1999/Cor 1:2001, *Technical Corrigendum to ISO/IEC 9899:1999*.

ISO/IEC 9899:1999/Cor 2:2004, *Technical Corrigendum to ISO/IEC 9899:1999*.

[ISO10026] ISO/IEC 10026-2:1998, *Information technology — Open Systems Interconnection — Distributed Transaction Processing — Part 2: OSI TP Service*.

[ISO10206] ISO/IEC 10206:1991, *Information technology — Programming languages — Extended Pascal*.

[ISO10646] ISO/IEC 10646:2003, *Information technology — Universal Multiple-Octet Coded Character Set (UCS)*.

[ISO11756] ISO/IEC 11756:1999, *Information technology — Programming languages — M*.

[ISO14651] ISO/IEC 14651:2007, *Information technology — International string ordering and comparison — Method for comparing character strings and description of the common template tailorable ordering*.

2.2 Other international standards

[IANA] The Internet Assigned Numbers Authority, *Character sets*
<http://www.iana.org/assignments/character-sets>

[Unicode] The Unicode Consortium, *The Unicode Standard*. (Information about the latest version of the Unicode standard can be found by using the "Latest Unicode Version" link on the "Enumerated Versions of The Unicode Standard" page.)
<http://www.unicode.org/versions/enumeratedversions.html>

[Unicode10] Davis, Mark and Whistler, Ken. *Unicode Technical Standard #10, Unicode Collation Algorithm, Version 4.1.0*, 2005-05-05. The Unicode Consortium.
<http://www.unicode.org/unicode/reports/tr10/tr10-14.html>

[Unicode15] Davis, Mark and Dürst, Martin, *Unicode Standard Annex #15: Unicode Normalization Forms, Version 4.1.0*, 2005-03-25, The Unicode Consortium
<http://www.unicode.org/reports/tr15/tr15-25.html>

[Unicode18] Davis, Mark, *Unicode Technical Standard #18: Unicode Regular Expressions*, The Unicode Consortium
<http://www.unicode.org/reports/tr18/>

[XQueryFO] (W3C Recommendation) XQuery 1.0 and XPath 2.0 Functions and Operators, 23 January 2007, World Wide Web Consortium.
<http://www.w3.org/TR/2007/REC-xpath-functions-20070123/>

3 Definitions, notations, and conventions

This Clause is modified by Clause 3, “Definitions, notations, and conventions”, in ISO/IEC 9075-3.
This Clause is modified by Clause 3, “Definitions, notations, and conventions”, in ISO/IEC 9075-4.
This Clause is modified by Clause 3, “Definitions, notations, and conventions”, in ISO/IEC 9075-10.
This Clause is modified by Clause 3, “Definitions, notations, and conventions”, in ISO/IEC 9075-11.
This Clause is modified by Clause 3, “Definitions, notations, and conventions”, in ISO/IEC 9075-13.
This Clause is modified by Clause 3, “Definitions, notations and conventions”, in ISO/IEC 9075-14.

3.1 Definitions

This Subclause is modified by Subclause 3.1, “Definitions”, in ISO/IEC 9075-3.
This Subclause is modified by Subclause 3.1, “Definitions”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 3.1, “Definitions”, in ISO/IEC 9075-10.
This Subclause is modified by Subclause 3.1, “Definitions”, in ISO/IEC 9075-13.
This Subclause is modified by Subclause 3.1, “Definitions”, in ISO/IEC 9075-14.

3.1.1 Definitions taken from [ISO10646]

For the purposes of this document, the definitions of the following terms given in [ISO10646] apply:

3.1.1.1 character

NOTE 1 — This is identical to *abstract character* as defined by [Unicode]. In ISO/IEC 9075, when the relevant character repertoire is UCS, a character can be thought of as *that which is represented by one code point*.

3.1.1.2 repertoire

3.1.2 Definitions taken from [ISO14651]

For the purposes of this document, the definitions of the following terms given in [14651] apply:

3.1.2.1 collation

3.1.3 Definitions taken from [Unicode]

For the purposes of this document, the definitions of the following terms given in [Unicode] apply:

3.1.3.1 character encoding form

3.1.3.2 code point

3.1 Definitions

3.1.3.3 code unit

3.1.3.4 control character

3.1.3.5 noncharacter

3.1.3.6 normalization

3.1.3.7 transcoding

3.1.4 Definitions taken from [ISO8601]

For the purposes of this document, the definitions of the following terms given in [ISO8601] apply:

3.1.4.1 Coordinated Universal Time (UTC)

3.1.4.2 date (date, calendar in [ISO8601])

3.1.5 Definitions taken from [ISO9075-1]

For the purposes of this document, the definitions of the following terms given in [ISO9075-1] apply.

3.1.5.1 atomic

3.1.5.2 compilation unit

3.1.5.3 data type

3.1.5.4 descriptor

3.1.5.5 fully qualified (of a name of some SQL object)

3.1.5.6 identifier

3.1.5.7 identify

3.1.5.8 implementation-defined

3.1.5.9 implementation-dependent

3.1.5.10 instance (of a value)

3.1.5.11 null value

3.1.5.12 object (as in 'x object')

3.1.5.13 persistent

3.1.5.14 property (of an object)

3.1.5.15 row

3.1.5.16 scope (of a standard)

3.1.5.17 scope (of a name or declaration)

- 3.1.5.18 **scope (of a reference type)**
- 3.1.5.19 **sequence**
- 3.1.5.20 **site**
- 3.1.5.21 **sizing item**
- 3.1.5.22 **SQL-connection**
- 3.1.5.23 **SQL-environment**
- 3.1.5.24 **SQL-implementation**
- 3.1.5.25 **SQL-session**
- 3.1.5.26 **SQL-statement**
- 3.1.5.27 **table**

3.1.6 Definitions provided in Part 2

For the purposes of this document, in addition to those definitions taken from other sources, the following definitions apply:

- 3.1.6.1 **assignable (of types, taken pairwise)**: The characteristic of a data type *T1* that permits a value of *T1* to be assigned to a site of a specified data type *T2*, where *T1* and *T2* may be the same data type.
- 3.1.6.2 **assignment**: The operation whose effect is to ensure that the value at a site *T* (known as the *target*) is identical to be a given value *S* (known as the *source*). Assignment is frequently indicated by the use of the phrase “*T* is set to *S*” or “the value of *T* is set to *S*”.
- 3.1.6.3 **attribute**: A component of a structured type. Each value *V* in structured type *T* has exactly one attribute value for each attribute *A* of *T*. The characteristics of an attribute are specified by an attribute descriptor. The value of an attribute may be retrieved as the result of the invocation *A(V)* of the observer function for that attribute.
- 3.1.6.4 **cardinality (of a collection)**: The number of elements in that collection. Those elements need not necessarily have distinct values. The objects to which this concept applies include tables and the values of collection types.
- 3.1.6.5 **comparable (of a pair of values)**: Capable of being compared, according to the rules of [Subclause 8.2](#), “<comparison predicate>”. In most, but not all, cases, the values of a data type can be compared one with another. For the specification of comparability of individual data types, see [Subclause 4.2](#), “Character strings”, through [Subclause 4.10](#), “Collection types”.
- 3.1.6.6 **constructor function**: A niladic SQL-invoked function of which exactly one is implicitly specified for every structured type. An invocation of the constructor function for data type *T* returns a value *V* of the most specific type of *T* such that *V* is not the null value and, for every observer function *O* defined for *T*, the invocation *O(V)* returns the default value of the attribute corresponding to *O*.
- 3.1.6.7 **declared type (of an expression denoting a value or anything that can be referenced to denote a value, such as, for example, a parameter, column, or variable)**: The unique data type that is common to every value that might result from evaluation of that expression.

3.1 Definitions

3.1.6.8 distinct (of a pair of comparable values): Capable of being distinguished within a given context.

NOTE 2 — Informally, two values are distinct if neither is null and the values are not equal. A null value and a non-null value are distinct. Two null values are not distinct. See Subclause 4.1.5, “Properties of distinct”, and the General Rules of Subclause 8.15, “<distinct predicate>”.

3.1.6.9 distinct type: A user-defined type derived from a predefined type.

3.1.6.10 duplicates: Two or more members of a multiset that are not distinct.

3.1.6.11 dyadic (of operators, functions, and procedures): Having exactly two operands or parameters.

NOTE 3 — An example of a dyadic operator in this part of ISO/IEC 9075 is “–”, specifying the subtraction of the right operand from the left operand. An example of a dyadic function is POSITION.

3.1.6.12 element type (of a collection type): The declared type specified in the definition of a collection type *CT*.

NOTE 4 — The declared type of every element of every value of type *CT* is *DT*.

3.1.6.13 equal (of a pair of comparable values): Yielding *True* if passed as arguments in a <comparison predicate> in which the <comp op> is <equals operator>. (see Subclause 8.2, “<comparison predicate>”).

3.1.6.14 external routine: An SQL-invoked routine whose routine body is an external body reference that identifies a program written in a programming language other than SQL.

3.1.6.15 fixed-length: A characteristic of character strings and binary strings that restricts a string to contain exactly one number of characters or octets, respectively, known as the length in characters or octets, respectively, of the string.

3.1.6.16 identical (of a pair of values): Indistinguishable, in the sense that it is impossible, by any means specified in ISO/IEC 9075, to detect any difference between them. For the full definition, see Subclause 9.8, “Determination of identical values”.

3.1.6.17 interface (of a user-defined type): The set comprising every function such that the declared type of at least one of its parameters or its result is that user-defined type.

3.1.6.18 monadic (of operators, functions, and procedures): Having exactly one operand or parameter.

NOTE 5 — An example of a monadic arithmetic operator in this part of ISO/IEC 9075 is “–”, specifying the negation of the operand. An example of a monadic function is CHARACTER_LENGTH, specifying the length in characters of the argument.

3.1.6.19 most specific type (of a value): The unique data type of which every data type of that value is a supertype.

3.1.6.20 mutator function: A dyadic, type-preserving SQL-invoked function implicitly defined by the definition of an attribute of a structured type that, when invoked, modifies the value of the attribute with which it is associated.

3.1.6.21 n-adic operator: An operator having a variable number of operands (informally: *n* operands).

NOTE 6 — An example of an *n*-adic operator in this part of ISO/IEC 9075 is COALESCE.

3.1.6.22 niladic (of functions and procedures): Having no parameters.

3.1.6.23 observer function: A monadic SQL-invoked function implicitly defined by the definition of an attribute of a structured type that, when invoked, returns the value of the attribute with which it is associated.

- 3.1.6.24 redundant duplicates:** All except one of any collection of duplicate values or rows.
- 3.1.6.25 REF value:** A value that references some site.
- 3.1.6.26 reference type:** A data type all of whose values are potential references to sites of one specified data type.
- 3.1.6.27 referenced type:** The declared type of the values at sites referenced by values of a particular reference type.
- 3.1.6.28 referenced value:** The value at the site referenced by a REF value.
- 3.1.6.29 result SQL parameter:** An SQL parameter that specifies RESULT.
- 3.1.6.30 result data type:** The declared type of the result of an SQL-invoked function.
- 3.1.6.31 result set:** A sequence of rows specified by a <cursor specification>, that is brought into existence by opening a cursor and ranged over by that cursor.
- 3.1.6.32 result set sequence:** A sequence of returned result sets.
- 3.1.6.33 returned result set:** A result set created during execution of an SQL-invoked procedure and not destroyed when that execution terminates. Such a result set can be accessed by using a cursor other than the one that brought it into existence (a received cursor).
- 3.1.6.34 savepoint:** A point within an SQL-transaction, identified by a savepoint name, to which that SQL-transaction may be restored.
- 3.1.6.35 signature (of an SQL-invoked routine):** The name of an SQL-invoked routine, the position and declared type of each of its SQL parameters, and an indication of whether it is an SQL-invoked function or an SQL-invoked procedure.
- 3.1.6.36 SQL argument:** An expression denoting a value to be substituted for an SQL parameter in an invocation of an SQL-invoked routine.
- 3.1.6.37 SQL-invoked routine:** A routine that is allowed to be invoked only from within SQL.
- 3.1.6.38 SQL parameter:** A parameter declared as part of the signature of an SQL-invoked routine.
- 3.1.6.39 SQL routine:** An SQL-invoked routine whose routine body is written in SQL.
- 3.1.6.40 structured type:** A user-defined type specified as a set of attributes.
- 3.1.6.41 subfield (of a row type):** A field that is a field of a row type *RT* or a field of a row type *RT2* that is the declared type of a field that is a subfield of *RT*.
- 3.1.6.42 subtype (of a data type):** A data type *T2* such that every value of *T2* is also a value of data type *T1*. If *T1* and *T2* are not compatible, then *T2* is a *proper subtype* of *T1*.
- NOTE 7 — “Compatible” is defined in Subclause 4.1.4, “Data type terminology”. See also **supertype**.
- 3.1.6.43 supertype (of a data type):** A data type *T1* such that every value of *T2* is also a value of data type *T1*. If *T1* and *T2* are not compatible, then *T1* is a *proper supertype* of *T2*.
- NOTE 8 — “Compatible” is defined in Subclause 4.1.4, “Data type terminology”. See also **subtype**.
- 3.1.6.44 transliteration:** A method of translating characters in one character set into characters of the same or a different character set.

3.1 Definitions

- 3.1.6.45 type-preserving function:** An SQL-invoked function, one of whose parameters is a result SQL parameter. The most specific type of the value returned by an invocation of a type-preserving function is identical to the most specific type of the SQL argument value substituted for the result SQL parameter.
- 3.1.6.46 user-defined type:** A type whose interface is user-defined.
- 3.1.6.47 variable-length:** A characteristic of character strings and binary strings that allows a string to contain any number of characters or octets, respectively, between 0 (zero) and some maximum number, known as the maximum length in characters or octets, respectively, of the string.
- 3.1.6.48 white space:** Characters used to separate tokens in SQL text; white space may be required (for example, to separate <nondelimiter token>s from one another) and may be used between any two tokens for which there are no rules prohibiting such use.

White space is any character in the Unicode General Category classes “Zs”, “Zl”, and “Zp”, or any of the following characters:

- U+0009, Horizontal Tabulation
- U+000A, Line Feed
- U+000B, Vertical Tabulation
- U+000C, Form Feed
- U+000D, Carriage Return
- U+0085, Next Line

NOTE 9 — The normative provisions of this International Standard impose no requirement that any character set have equivalents for any of these characters except U+0020 (<space>); however, by reference to this definition of white space, they do impose the requirement that every equivalent for one of these shall be recognized as a white space character.

NOTE 10 — The Unicode General Category classes “Zs”, “Zl”, and “Zp” are assigned to Unicode characters that are, respectively, space separators, line separators, and paragraph separators.

The only character that is a member of the Unicode General Category class “Zl” is U+2028 (Line Separator). The only character that is a member of the Unicode General Category class “Zp” is U+2029 (Paragraph Separator). The characters that are members of the Unicode General Category class “Zs” are: U+0020 (Space), U+00A0 (No-Break Space), U+1680 (Ogham Space Mark), U+180E (Mongolian Vowel Separator), U+2000 (En Quad), U+2001 (Em Quad), U+2002 (En Space), U+2003 (Em Space), U+2004 (Three-Per-Em Space), U+2005 (Four-Per-Em Space), U+2006 (Six-Per-Em Space), U+2007 (Figure Space), U+2008 (Punctuation Space), U+2009 (Thin Space), U+200A (Hair Space), U+202F (Narrow No-Break Space), U+205F (Space, Medium Mathematical), and U+3000 (Ideographic Space).

- 3.1.6.49 with-return cursor:** A cursor that, when opened, creates a result set that is capable of becoming a returned result set. The WITH RETURN option of <declare cursor>, <dynamic declare cursor>, and <allocate cursor statement> specifies a with-return cursor. WITH RETURN may also be specified in the content of an <attributes variable> in a <prepare statement>, to indicate that the prepared statement, when opened as a dynamic cursor, creates a with-return cursor.
- 3.1.6.50 XQuery captured substring:** a captured substring, as defined by [XQueryFO].
- 3.1.6.51 XQuery option flag:** a valid value of the \$flags argument of fn:matches, as defined by [XQueryFO].
- 3.1.6.52 XQuery regular expression:** a regular expression, as defined by [XQueryFO].

3.1.6.53 XQuery regular expression parenthesized sub-expression: a regular expression parenthesized sub-expression, as defined by [XQueryFO].

3.1.6.54 XQuery replacement string: a valid value of the \$replacement argument of fn:replace, as defined by [XQueryFO].

3.2 Notation

This Subclause is modified by Subclause 3.2, “Notation”, in ISO/IEC 9075-14.

The notation used in this document is defined in [ISO9075-1].

3.3 Conventions

This Subclause is modified by Subclause 3.2, “Conventions”, in ISO/IEC 9075-3.

This Subclause is modified by Subclause 3.1, “Conventions”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 3.2, “Conventions”, in ISO/IEC 9075-10.

This Subclause is modified by Subclause 3.1, “Conventions”, in ISO/IEC 9075-11.

This Subclause is modified by Subclause 3.2, “Conventions”, in ISO/IEC 9075-13.

The conventions used in this document are defined in [ISO9075-1], with the following additions.

3.3.1 Use of terms

This Subclause is modified by Subclause 3.1.1, “Use of terms”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 3.2.1, “Use of terms”, in ISO/IEC 9075-10.

3.3.1.1 Other terms

This Subclause is modified by Subclause 3.1.1.1, “Other terms”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 3.2.1.1, “Other terms”, in ISO/IEC 9075-10.

An SQL-statement *SI* is said to be executed as a *direct result of executing an SQL-statement* if *SI* is the SQL-statement contained in an <externally-invoked procedure> or <SQL-invoked routine> that has been executed.

An SQL-statement *SI* is said to be executed as a *direct result of executing an SQL-statement* if *SI* is the value of an <SQL statement variable> referenced by an <execute immediate statement> contained in an <externally-invoked procedure> that has been executed, or if *SI* was the value of the <SQL statement variable> that was associated with an <SQL statement name> by a <prepare statement> and that same <SQL statement name> is referenced by an <execute statement> contained in an <externally-invoked procedure> that has been executed.

(Blank page)

4 Concepts

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-3.

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-4.

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-9.

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-10.

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-11.

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-13.

This Clause is modified by [Clause 4, “Concepts”](#), in ISO/IEC 9075-14.

4.1 Data types

This Subclause is modified by [Subclause 4.1, “Data types”](#), in ISO/IEC 9075-9.

This Subclause is modified by [Subclause 4.7, “Data types”](#), in ISO/IEC 9075-13.

This Subclause is modified by [Subclause 4.1, “Data types”](#), in ISO/IEC 9075-14.

4.1.1 General introduction to data types

A data type is a set of representable values. Every representable value belongs to at least one data type and some belong to several data types.

Exactly one of the data types of a value *V*, namely the most specific type of *V*, is a subtype of every data type of *V*. A <value expression> *E* has exactly one declared type, common to every possible result of evaluating *E*. Items that can be referenced by name, such as SQL parameters, columns, fields, attributes, and variables, also have declared types.

SQL supports three sorts of data types: *predefined data types*, *constructed types*, and *user-defined types*. Predefined data types are sometimes called “built-in data types”, though not in this International Standard. User-defined types can be defined by a standard, by an implementation, or by an application.

A constructed type is specified using one of SQL's data type constructors, ARRAY, MULTISSET, REF, and ROW. A constructed type is either an array type, a multiset type, a reference type, or a row type, according to whether it is specified with ARRAY, MULTISSET, REF, or ROW, respectively. Array types and multiset types are known generically as *collection types*.

Every predefined data type is a subtype of itself and of no other data types. It follows that every predefined data type is a supertype of itself and of no other data types. The predefined data types are individually described in each of [Subclause 4.2, “Character strings”](#), through [Subclause 4.6, “Datetimes and intervals”](#).

Row types, reference types, and collection types are described in [Subclause 4.8, “Row types”](#), [Subclause 4.9, “Reference types”](#), [Subclause 4.10, “Collection types”](#), respectively.

A user-defined type is either a *distinct type* or a *structured type*. User-defined types are described in [Subclause 4.7, “User-defined types”](#).

4.1 Data types

4.1.2 Naming of predefined types

This Subclause is modified by Subclause 4.1.1, “Naming of predefined types”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 4.1.1, “Naming of predefined types”, in ISO/IEC 9075-14.

09 14 SQL defines predefined data types named by the following <key word>s:

CHARACTER

CHARACTER VARYING

CHARACTER LARGE OBJECT

BINARY

BINARY VARYING

BINARY LARGE OBJECT

NUMERIC

DECIMAL

SMALLINT

INTEGER

BIGINT

FLOAT

REAL

DOUBLE PRECISION

BOOLEAN

DATE

TIME

TIMESTAMP

INTERVAL

These names are used in the *type designators* that constitute the *type precedence lists* specified in Subclause 9.5, “Type precedence list determination”.

09 14 For reference purposes:

- The data types CHARACTER, CHARACTER VARYING, and CHARACTER LARGE OBJECT are collectively referred to as *character string types* and the values of character string types are known as *character strings*.
- The data types BINARY, BINARY VARYING, and BINARY LARGE OBJECT are referred to as *binary string types* and the values of binary string types are referred to as *binary strings*.

- The data types CHARACTER LARGE OBJECT and BINARY LARGE OBJECT are collectively referred to as *large object string types* and the values of large object string types are referred to as *large object strings*.
- Character string types and binary string types are collectively referred to as *string types* and values of string types are referred to as *strings*.
- The data types NUMERIC, DECIMAL, SMALLINT, INTEGER, and BIGINT are collectively referred to as *exact numeric types*.
- The data types FLOAT, REAL, and DOUBLE PRECISION are collectively referred to as *approximate numeric types*.
- Exact numeric types and approximate numeric types are collectively referred to as *numeric types*. Values of numeric types are referred to as *numbers*.
- The data types TIME WITHOUT TIME ZONE and TIME WITH TIME ZONE are collectively referred to as *time types* (or, for emphasis, as time with or without time zone).
- The data types TIMESTAMP WITHOUT TIME ZONE and TIMESTAMP WITH TIME ZONE are collectively referred to as *timestamp types* (or, for emphasis, as timestamp with or without time zone).
- The data types DATE, TIME, and TIMESTAMP are collectively referred to as *datetime types*.
- Values of datetime types are referred to as *datetimes*.
- The data type INTERVAL is referred to as an *interval type*. Values of interval types are called *intervals*.

Each data type has an associated data type descriptor; the contents of a data type descriptor are determined by the specific data type that it describes. A data type descriptor includes an identification of the data type and all information needed to characterize a value of that data type.

Subclause 6.1, “<data type>”, describes the semantic properties of each data type.

4.1.3 Host language data types

This Subclause is modified by Subclause 4.7.1, “Host language data types”, in ISO/IEC 9075-13.

13 Each host language has its own data types, which are separate and distinct from SQL data types, even though similar names may be used to describe the data types. Mappings of SQL data types to data types in host languages are described in Subclause 11.51, “<SQL-invoked routine>”, and Subclause 21.1, “<embedded SQL host program>”. Not every SQL data type has a corresponding data type in every host language.

4.1.4 Data type terminology

This Subclause is modified by Subclause 4.1.2, “Data type terminology”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 4.1.2, “Data type terminology”, in ISO/IEC 9075-14.

The notion of a *constituent* of a declared type *DT* is defined recursively as follows:

- *DT* is a constituent of *DT*.

4.1 Data types

- If *DT* is a row type, then the declared type of each field of *DT* is a constituent of *DT*.
- If *DT* is a collection type, then the element type of *DT* is a constituent of *DT*.
- Every constituent of a constituent of *DT* is a constituent of *DT*.

A data type *TY* is *usage-dependent* on a user-defined type *UDT* if one of the following conditions is true:

- *TY* is *UDT*.
- *TY* is a reference type whose referenced type is *UDT*.
- *TY* is a row type, and the declared type of some field of *TY* is usage-dependent on *UDT*.
- *TY* is a collection type, and the declared element type of *TY* is usage-dependent on *UDT*.

A structured type *ST* is *directly based on* a data type *DT* if any of the following are true:

- *DT* is the declared type of some attribute of *ST*.
- *DT* is a direct supertype of *ST*.
- *DT* is a direct subtype of *ST*.
- *DT* is compatible with *ST*.

A collection type *CT* is *directly based on* a data type *DT* if *DT* is the element type of *CT*.

A row type *RT* is *directly based on* a data type *DT* if *DT* is the declared type of some field (or the data type of the domain of some field) whose descriptor is included in the descriptor of *RT*.

A data type *DT1* is *based on* a data type *DT2* if *DT1* is compatible with *DT2*, *DT1* is directly based on *DT2*, or *DT1* is directly based on some data type that is based on *DT2*.

Two data types, *T1* and *T2*, are said to be *compatible* if *T1* is assignable to *T2*, *T2* is assignable to *T1*, and their descriptors include the same data type name. If they are row types, it shall further be the case that the declared types of their corresponding fields are pairwise compatible. If they are collection types, it shall further be the case that their element types are compatible. If they are reference types, it shall further be the case that their referenced types are compatible.

NOTE 11 — The data types “CHARACTER(*n*) CHARACTER SET *CS1*” and “CHARACTER(*m*) CHARACTER SET *CS2*”, where *CS1* ≠ *CS2*, have descriptors that include the same data type name (CHARACTER), but are not mutually assignable; therefore, they are not compatible.

Ordering and comparison of values of the predefined data types requires knowledge only about those predefined data types. However, to be able to compare and order values of constructed types or of user-defined types, additional information is required. We say that some type *T* is *S-ordered*, for some set of types *S*, if, in order to compare and order values of type *T*, information about ordering at least one of the types in *S* is first required. A definition of *S-ordered* is required for several *S* (that is, for several sets of types), but not for all possible such sets.

The general definition of *S-ordered* is as follows:

Let *T* be a type and let *S* be a set of types. *T* is *S-ordered* if one of the following is true:

- *T* is a member of *S*.
- *T* is a row type and the declared type of some field of *T* is *S-ordered*.
- *T* is a collection type and the element type of *T* is *S-ordered*.

- *T* is a structured type whose comparison form is STATE and the declared type of some attribute of *T* is *S*-ordered.
- *T* is a user-defined type whose comparison form is MAP and the return type of the SQL-invoked function that is identified by the <map function specification> is *S*-ordered.
- *T* is a reference type with a derived representation and the declared type of some attribute enumerated by the <derived representation> is *S*-ordered.

09 14 The notion of *S*-ordered is applied in the following definitions:

- A type *T* is *LOB-ordered* if *T* is *S*-ordered, where *S* is the set of large object types.
- A type *T* is *array-ordered* if *T* is *S*-ordered, where *S* is the set of array types.
- A type *T* is *multiset-ordered* if *T* is *S*-ordered, where *S* is the set of multiset types.
- A type *T* is *reference-ordered* if *T* is *S*-ordered, where *S* is the set of reference types.
- A type *T* is *DT-EC-ordered* if *T* is *S*-ordered, where *S* is the set of distinct types with EQUALS ONLY comparison form (DT-EC stands for “distinct type-equality comparison”).
- A type *T* is *DT-FC-ordered* if *T* is *S*-ordered, where *S* is the set of distinct types with FULL comparison form.
- A type *T* is *DT-NC-ordered* if *T* is *S*-ordered, where *S* is the set of distinct types with no comparison form.
- A type *T* is *ST-EC-ordered* if *T* is *S*-ordered, where *S* is the set of structured types with EQUALS ONLY comparison form.
- A type *T* is *ST-FC-ordered* if *T* is *S*-ordered, where *S* is the set of structured types with FULL comparison form.
- A type *T* is *ST-NC-ordered* if *T* is *S*-ordered, where *S* is the set of structured types with no comparison form.
- A type *T* is *ST-ordered* if *T* is ST-EC-ordered, ST-FC-ordered, or ST-NC-ordered.
- A type *T* is *UDT-EC-ordered* if *T* is either DT-EC-ordered or ST-EC-ordered (UDT stands for “user-defined type”).
- A type *T* is *UDT-FC-ordered* if *T* is either DT-FC-ordered or ST-FC-ordered
- A type *T* is *UDT-NC-ordered* if *T* is either DT-NC-ordered or ST-NC-ordered.

4.1.5 Properties of distinct

Two comparable values are distinct if they are capable of being distinguished within a given context.

Two null values are not distinct.

A null value and a non-null value are distinct.

Two non-null values are distinct if the General Rules of [Subclause 8.15](#), “<distinct predicate>”, return *True*.

4.1 Data types

The result of evaluating whether or not two comparable values are distinct is never *Unknown*. The result of evaluating whether or not two values that are not comparable (for example, values of a user-defined type that has no comparison type) are distinct is not defined.

4.2 Character strings

This Subclause is modified by Subclause 4.2, “Character strings”, in ISO/IEC 9075-10.

4.2.1 Introduction to character strings

A character string is a sequence of characters. All the characters in a character string are taken from a single character set. A character string has a length, which is the number of characters in the sequence. The length is 0 (zero) or a positive integer.

A *character string type* is described by a character string type descriptor. A character string type descriptor contains:

- The name of the specific character string type (CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT; NATIONAL CHARACTER, NATIONAL CHARACTER VARYING, and NATIONAL CHARACTER LARGE OBJECT are represented as CHARACTER, CHARACTER VARYING, and CHARACTER LARGE OBJECT, respectively).
- The length or maximum length in characters of the character string type.
- The catalog name, schema name, and character set name of the character set of the character string type.
- The catalog name, schema name, and collation name of the collation of the character string type.

A *character large object type* is a character string type where the name of the specific character string type is CHARACTER LARGE OBJECT. A value of a character large object type is a *large object character string*.

The character set of a character string type may be specified explicitly or implicitly.

The <key word>s NATIONAL CHARACTER are used to specify an implementation-defined character set. Special syntax (N' *string* ') is provided for representing literals in that character set.

With two exceptions, a character string expression is assignable only to sites of a character string type whose character set is the same. The exceptions are as specified in Subclause 4.2.8, “Universal character sets”, and such other cases as may be implementation-defined. If a store assignment would result in the loss of non-*<space>* characters due to truncation, then an exception condition is raised. If a retrieval assignment or evaluation of a *<cast specification>* would result in the loss of characters due to truncation, then a warning condition is raised.

Character sets fall into three categories: those defined by national or international standards, those defined by SQL-implementations, and those defined by applications. The character sets defined by [ISO10646] and [Unicode] are known as *Universal Character Sets* (UCS) and their treatment is described in Subclause 4.2.8, “Universal character sets”. Every character set contains the *<space>* character (equivalent to U+0020). An application defines a character set by assigning a new name to a character set from one of the first two categories. They can be defined to “reside” in any schema chosen by the application. Character sets defined by standards or by SQL-implementations reside in the Information Schema (named INFORMATION_SCHEMA) in each

catalog, as do collations defined by standards and collations, transliterations, and transcodings defined by SQL-implementations.

NOTE 12 — The Information Schema is defined in [ISO9075-11].

4.2.2 Comparison of character strings

Two character strings are comparable if and only if either they have the same character set or there exists at least one collation that is applicable to both their respective character sets.

A *collation* is defined by [ISO14651] as “a process by which two strings are determined to be in exactly one of the relationships of less than, greater than, or equal to one another”. Each collation known in an SQL-environment is applicable to one or more character sets, and for each character set, one or more collations are applicable to it, one of which is associated with it as its *character set collation*.

Anything that has a declared type can, if that type is a character string type, be associated with a collation applicable to its character set; this is known as a *declared type collation*. Every declared type that is a character string type has a collation derivation, this being either *none*, *implicit*, or *explicit*. The collation derivation of a declared type with a declared type collation that is explicitly or implicitly specified by a <data type> is *implicit*. If the collation derivation of a declared type that has a declared type collation is not *implicit*, then it is *explicit*. The collation derivation of an expression of character string type that has no declared type collation is *none*.

An operation that explicitly or implicitly involves character string comparison is a *character comparison operation*. At least one of the operands of a character comparison operation shall have a declared type collation.

There may be an SQL-session collation for some or all of the character sets known to the SQL-implementation (see [Subclause 4.37, “SQL-sessions”](#)).

The collation used for a particular character comparison is specified by [Subclause 9.13, “Collation determination”](#).

The comparison of two character string expressions depends on the collation used for the comparison (see [Subclause 9.13, “Collation determination”](#)). When values of unequal length are compared, if the collation for the comparison has the NO PAD characteristic and the shorter value is equal to some prefix of the longer value, then the shorter value is considered less than the longer value. If the collation for the comparison has the PAD SPACE characteristic, for the purposes of the comparison, the shorter value is effectively extended to the length of the longer by concatenation of <space>s on the right.

For every character set, there is at least one collation.

4.2.3 Operations involving character strings

4.2.3.1 Regular expression syntaxes

This International Standard utilizes two syntaxes for regular expressions, a POSIX-based regular expression syntax and the XQuery regular expression syntax.

The POSIX-based regular expression syntax is similar to the syntax specified in [POSIX], but is normatively defined in the General Rules of [Subclause 8.6, “<similar predicate>”](#).

The XQuery regular expression syntax is normatively defined in [XQueryFO].

4.2.3.2 Operators that operate on character strings and return character strings

<concatenation operator> is an operator, `|`, that returns the character string made by joining its character string operands in the order given.

<character substring function> is a triadic function, SUBSTRING, that returns a string extracted from a given string according to a given numeric starting position and a given numeric length.

<regular expression substring function> is a triadic function, SUBSTRING, distinguished by the keywords SIMILAR and UESCAPE. It has three parameters: a source character string, a pattern string, and an escape character. It returns a result string extracted from the source character string by pattern matching using POSIX-based regular expressions.

- **Step 1:** The escape character is exactly one character in length. As indicated in Figure 1, “Operation of <regular expression substring function>”, the escape character precedes two instances of <double quote> that are used to partition the pattern string into three subpatterns (identified as *R1*, *R2*, and *R3*).
- **Step 2:** If the source string *S* does not satisfy the predicate

```
'S' SIMILAR TO 'R1' || 'R2' || 'R3'
```

then the result is the null value.

- **Step 3:** Otherwise, *S* is partitioned into two substrings *S1* and *S23* such that *S1* is the shortest initial substring of *S* such that the following condition is satisfied:

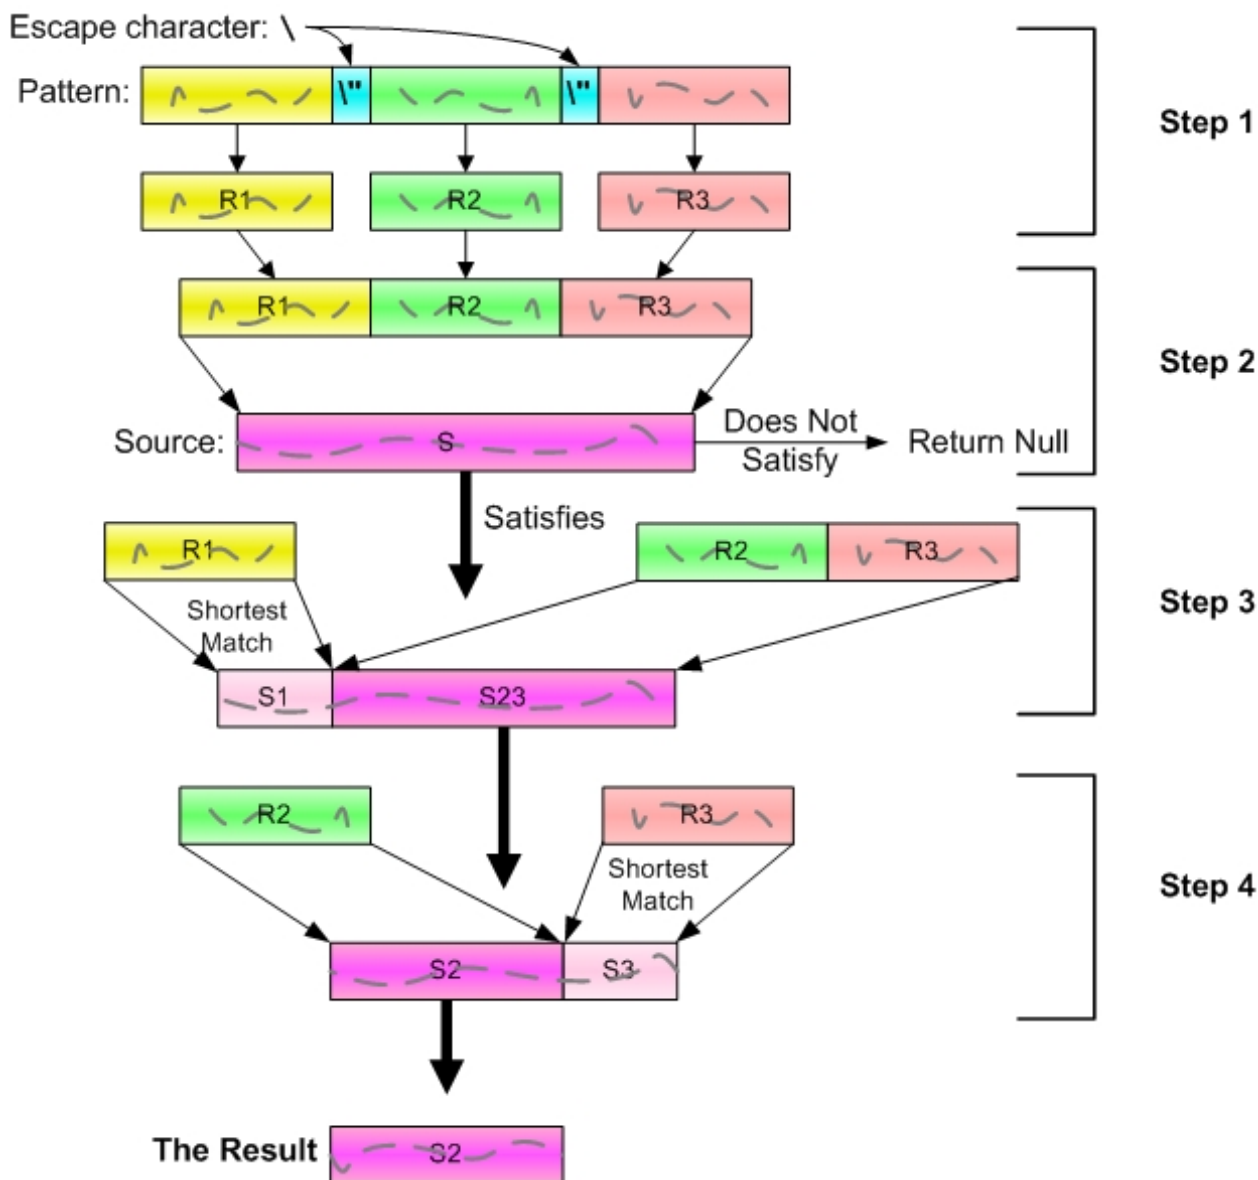
```
'S1' SIMILAR TO 'R1' AND  
'S23' SIMILAR TO '(' || 'R2' || 'R3' || ')'
```

- **Step 4:** Next, *S23* is partitioned into two substrings *S2* and *S3* such that *S3* is the shortest final substring such that the following condition is satisfied:

```
'S2' SIMILAR TO 'R2' AND 'S3' SIMILAR TO 'R3'
```

The result of the <regular expression substring function> is *S2*.

Figure 1 — Operation of <regular expression substring function>



<regex substring function> is a function, SUBSTRING_REGEX, that searches a string for an XQuery regular expression pattern and returns one occurrence of the matching substring.

<regex transliteration> is a function, TRANSLATE_REGEX, that searches a string for an XQuery regular expression pattern and returns the string with either one or every occurrence of the XQuery regular expression pattern replaced using an XQuery replacement string.

<character overlay function> is a function, OVERLAY, that modifies a string argument by replacing a given substring of the string, which is specified by a given numeric starting position and a given numeric length, with another string (called the replacement string). When the length of the substring is zero, nothing is removed

from the original string and the string returned by the function is the result of inserting the replacement string into the original string at the starting position.

<fold> is a pair of functions for converting all the lower case and title case characters in a given string to upper case (UPPER) or all the upper case and title case characters to lower case (LOWER). A lower case character is a character in the Unicode General Category class “Ll” (lower-case letters). An upper case character is a character in the Unicode General Category class “Lu” (upper-case letters). A title case character is a character in the Unicode General Category class “Lt” (title-case letters).

NOTE 13 — Case correspondences are not always one-to-one: the result of case folding may be of a different length in characters than the source string. For example, U+00DF, “ß”, Latin Small Letter Sharp S, becomes “SS” when folded to upper case.

<transcoding> is a function that invokes an installation-supplied transcoding to return a character string *S2* derived from a given character string *S1*. It is intended, though not enforced by this part of ISO/IEC 9075, that *S2* be exactly the same sequence of characters as *S1*, but encoded according to some different character encoding form. A typical use might be to convert a character string from two-octet UCS to one-octet Latin1 or *vice versa*.

<trim function> is a function that returns its first string argument with leading and/or trailing pad characters removed. The second argument indicates whether leading, or trailing, or both leading and trailing pad characters should be removed. The third argument specifies the pad character that is to be removed.

<character transliteration> is a function for changing each character of a given string according to some many-to-one or one-to-one mapping between two not necessarily distinct character sets. The mapping, rather than being specified as part of the function, is some external function identified by a <transliteration name>.

For any pair of character sets, there are zero or more transliterations that may be invoked by a <character transliteration>. A transliteration is described by a transliteration descriptor. A transliteration descriptor includes:

- The name of the transliteration.
- The name of the character set from which it translates.
- The name of the character set to which it translates.
- The specific name of the SQL-invoked routine that performs the transliteration.

4.2.3.3 Other operators involving character strings

<length expression> returns the length of a given character string, as an exact numeric value, in characters or octets according to the choice of function.

<position expression> determines the first position, if any, at which one string, *S1*, occurs within another, *S2*. If *S1* is of length zero, then it occurs at position 1 (one) for any value of *S2*. If *S1* does not occur in *S2*, then zero is returned. The declared type of a <position expression> is exact numeric.

<like predicate> uses the triadic operator LIKE (or the inverse, NOT LIKE), operating on three character strings and returning a Boolean. LIKE determines whether or not a character string “matches” a given “pattern” (also a character string). The characters <percent> and <underscore> have special meaning when they occur in the pattern. The optional third argument is a character string containing exactly one character, known as the “escape character”, for use when a <percent>, <underscore>, or the “escape character” itself is required in the pattern without its special meaning.

<similar predicate> uses the triadic operator SIMILAR (or the inverse, NOT SIMILAR), operating on three character strings and returning a Boolean. SIMILAR determines whether or not a character string “matches” a given “pattern” (also a character string) using a POSIX-based regular expression. The pattern is in the form of a “regular expression”. In this regular expression, certain characters (<left bracket>, <right bracket>, <left paren>, <right paren>, <vertical bar>, <circumflex>, <minus sign>, <plus sign>, <asterisk>, <underscore>, <percent>, <question mark>, <left brace>) have a special meaning. The optional third argument specifies the “escape character”, for use when one of the special characters or the “escape character” itself is required in the pattern without its special meaning.

<regex position expression> is a function, POSITION_REGEX, that searches a string for an XQuery regular expression pattern and returns an integer indicating the beginning position, or 1 (one) plus the ending position, of one occurrence of the matched substring.

<regex like predicate> is a predicate, LIKE_REGEX, that performs XQuery regular expression matching.

4.2.3.4 Operations involving large object character strings

Large object character strings cannot be operated on by all string operations. Large object character strings can, however, be operated on by the following operations:

- <null predicate>.
- <like predicate>.
- <similar predicate>.
- <position expression>.
- <comparison predicate> with an <equals operator> or <not equals operator>.
- <quantified comparison predicate> with the <equals operator> or <not equals operator>.

As a result of these restrictions, large object character strings cannot be used in (among other places):

- predicates other than those listed above and the <exists predicate>
- <general set function>.
- <group by clause>.
- <order by clause>.
- <unique constraint definition>.
- <referential constraint definition>.
- <select list> of a <query specification> that has a <set quantifier> of DISTINCT.
- UNION, INTERSECT, and EXCEPT.
- columns used for matching when forming a <joined table>.

All the operations described within Subclause 4.2.3.2, “Operators that operate on character strings and return character strings”, and Subclause 4.2.3.3, “Other operators involving character strings”, are supported for large object character strings.

4.2.4 Character repertoires

An SQL-implementation supports one or more character repertoires. These character repertoires may be defined by a standard or be implementation-defined.

A character repertoire is described by a character repertoire descriptor. A character repertoire descriptor includes:

- The name of the character repertoire.
- The name of the default collation for the character repertoire.

The following character repertoire names are specified as part of ISO/IEC 9075:

- SQL_CHARACTER is a character repertoire that consists of the 88 <SQL language character>s as specified in [Subclause 5.1](#), “<SQL terminal character>”. The name of the default collation is SQL_CHARACTER.
- GRAPHIC_IRV is the character repertoire that consists of the 95-character graphic subset of the International Reference Version (IRV) as specified in [ISO646]. Its repertoire is a proper superset of that of SQL_CHARACTER. The name of the default collation is GRAPHIC_IRV.
- LATIN1 is the character repertoire defined in [ISO8859-1]. The name of the default collation is LATIN1.
- ISO8BIT is the character repertoires formed by combining the character repertoire specified by [ISO8859-1] and the “control characters” specified by [ISO6429]. The repertoire consists of all 255 characters, each consisting of exactly 8 bits, as, including all control characters and all graphic characters except the character corresponding to the numeric value 0 (zero). The name of the default collation is ISO8BIT.
- UCS is the Universal Character Set repertoire specified by [Unicode] and by [ISO10646]. It is implementation-defined whether the name of the default collation is UCS_BASIC or UNICODE.
- SQL_TEXT is a character repertoire that is an implementation-defined subset of the repertoire of the Universal Character Set that includes every <SQL language character> and every character in every character set supported by the SQL-implementation. The name of the default collation is SQL_TEXT.
- SQL_IDENTIFIER is a character repertoire consisting of the <SQL language character>s and all other characters that the SQL-implementation supports for use in <regular identifier>s. The name of the default collation is SQL_IDENTIFIER.

4.2.5 Character encoding forms

An SQL-implementation supports one or more character encoding forms for each character repertoire that it supports. These character encoding forms may be defined by a standard or be implementation-defined.

A character encoding form is described by a character encoding form descriptor. A character encoding form descriptor includes:

- The name of the character encoding form.
- The name of the character repertoire to which it is applicable.

The following character encoding form names are specified as part of ISO/IEC 9075:

- SQL_CHARACTER is an implementation-defined character encoding form. It is applicable to the SQL_CHARACTER character repertoire.

- GRAPHIC_IRV is the character encoding form in which the coded representation of each character is specified in [ISO646]. It is applicable to the GRAPHIC_IRV character repertoire.
- LATIN1 is the character encoding form specified in [ISO8859-1]. It is applicable to the LATIN1 character repertoire.
- ISO8BIT is the character encoding form specified in [ISO8859-1], augmented by [ISO6429]. When restricted to the LATIN1 characters, it is the same character encoding form as LATIN1. It is applicable to the ISO8BIT character repertoire.
- UTF32 is the character encoding form specified in [ISO10646], in which each character is encoded as four octets. It is applicable to the UCS character repertoire.
- UTF16 is the character encoding form specified in [ISO10646], in which each character is encoded as two or four octets. It is applicable to the UCS character repertoire.
- UTF8 is the character encoding form specified in [ISO10646], in which each character is encoded as from one to four octets. It is applicable to the UCS character repertoire.
- SQL_TEXT is an implementation-defined character encoding form. It is applicable to the SQL_TEXT character repertoire.
- SQL_IDENTIFIER is an implementation-defined character encoding form. It is applicable to the SQL_IDENTIFIER character repertoire.

If an SQL-implementation supplies more than one character encoding form for a particular character repertoire, then it shall specify a precedence ordering of the character encoding forms of that character repertoire. The precedence of character encoding forms applicable to the UCS character repertoire and defined in this part of ISO/IEC 9075 is:

UTF8 < UTF16 < UTF32

4.2.6 Collations

An SQL-implementation supports one or more collations for each character repertoire that it supports, and one or more collations for each character set that it supports. A collation is described by a collation descriptor. A collation descriptor includes:

- The name of the collation.
- The name of the character repertoire to which it is applicable.
- A list of the names of the character sets to which the collation can be applied.
- Whether the collation has the NO PAD or the PAD SPACE characteristic.

The supported collation names are specified as part of ISO/IEC 9075:

- SQL_CHARACTER is an implementation-defined collation. It is applicable to the SQL_CHARACTER character repertoire.
- GRAPHIC_IRV is a collation in which the ordering is determined by treating the code points defined by [ISO646] as unsigned integers. It is applicable to the GRAPHIC_IRV character repertoire.

4.2 Character strings

- LATIN1 is a collation in which the ordering is determined by treating the code points defined by [ISO8859-1] as unsigned integers. It is applicable to the LATIN1 character repertoire.
- ISO8BIT is a collation in which the ordering is determined by treating the code points defined by [ISO8859-1] as unsigned integers. When restricted to the LATIN1 characters, it produces the same collation as LATIN1. It is applicable to the ISO8BIT character repertoire.
- UCS_BASIC is a collation in which the ordering is determined entirely by the Unicode scalar values of the characters in the strings being sorted. It is applicable to the UCS character repertoire. Since every character repertoire is a subset of the UCS repertoire, the UCS_BASIC collation is potentially applicable to every character set.

NOTE 14 — The Unicode scalar value of a character is its code point treated as an unsigned integer.

- UNICODE is the collation in which the ordering is determined by applying the Unicode Collation Algorithm with the Default Unicode Collation Element Table, as specified in [Unicode10]. It is applicable to the UCS character repertoire. Since every character repertoire is a subset of the UCS repertoire, the UNICODE collation is potentially applicable to every character set.
- SQL_TEXT is an implementation-defined collation. It is applicable to the SQL_TEXT character repertoire.
- SQL_IDENTIFIER is an implementation-defined collation. It is applicable to the SQL_IDENTIFIER character repertoire.

4.2.7 Character sets

This Subclause is modified by Subclause 4.2.2, “Character sets”, in ISO/IEC 9075-10.

An SQL <character set specification> allows a reference to a character set name defined by a standard, an SQL-implementation, or a user.

A character set is described by a character set descriptor. A character set descriptor includes:

- The name of the character set.
- The name of the character repertoire for the character set.
- The name of the character encoding form for the character set.
- The name of the default collation for the character set.

10 The following SQL supported character set names are specified as part of ISO/IEC 9075:

- 10 SQL_CHARACTER is a character set whose repertoire is SQL_CHARACTER and whose character encoding form is SQL_CHARACTER. The name of its default collation is SQL_CHARACTER.
- GRAPHIC_IRV is a character set whose repertoire is GRAPHIC_IRV and whose character encoding form is GRAPHIC_IRV. The name of its default collation is GRAPHIC_IRV.
- ASCII_GRAPHIC is a synonym for GRAPHIC_IRV.
- LATIN1 is a character set whose repertoire is LATIN1 and whose character encoding form is LATIN1. The name of its default collation is LATIN1.

- ISO8BIT is a character set whose repertoire is ISO8BIT and whose character encoding form is ISO8BIT. The name of its default collation is ISO8BIT.
- ASCII_FULL is a synonym for ISO8BIT.
- UTF32 is a character set whose repertoire is UCS and whose character encoding form is UTF32. It is implementation-defined whether the name of its default collation is UCS_BASIC or UNICODE.
- UTF16 is a character set whose repertoire is UCS and whose character encoding form is UTF16. It is implementation-defined whether the name of its default collation is UCS_BASIC or UNICODE.
- UTF8 is the name of a character set whose repertoire is UCS and whose character encoding form is UTF8. It is implementation-defined whether the name of its default collation is UCS_BASIC or UNICODE.
- SQL_TEXT is a character set whose repertoire is SQL_TEXT and whose character encoding form is SQL_TEXT. The name of its default collation is SQL_TEXT.
- SQL_IDENTIFIER is a character set whose repertoire is SQL_IDENTIFIER and whose character encoding form is SQL_IDENTIFIER. The name of its default collation is SQL_IDENTIFIER.

The result of evaluating a character string expression whose most specific type has character set *CS* is constrained to consist of characters drawn from the character repertoire of *CS*.

Table 1 — Overview of character sets

Character Set	Character Repertoire	Character Encoding Form	Collation	Synonym
GRAPHIC_IRV	GRAPHIC_IRV	GRAPHIC_IRV	GRAPHIC_IRV	ASCII_GRAPHIC
ISO8BIT	ISO8BIT	ISO8BIT	ISO8BIT	ASCII_FULL
LATIN1	LATIN1	LATIN1	LATIN1	
SQL_CHARACTER	SQL_CHARACTER	SQL_CHARACTER	SQL_CHARACTER	
SQL_TEXT	SQL_TEXT	SQL_TEXT	SQL_TEXT	
SQL-IDENTIFIER	SQL-IDENTIFIER	SQL-IDENTIFIER	SQL-IDENTIFIER	
UTF16	UCS	UTF16	UCS_BASIC or UNICODE	
UTF32	UCS	UTF32	UCS_BASIC or UNICODE	
UTF8	UCS	UTF8	UCS_BASIC or UNICODE	

NOTE 15 — An SQL-implementation may supply additional character sets and/or additional character encoding forms and collations for character sets defined in this Part of ISO/IEC 9075.

4.2.8 Universal character sets

A *UCS string* is a character string whose character repertoire is UCS and whose character encoding form is one of UTF8, UTF16, or UTF32. Any two UCS strings are comparable.

An SQL-implementation may assume that all UCS strings are normalized in one of Normalization Form C (NFC), Normalization Form D (NFD), Normalization Form KC (NFKC), or Normalization Form KD (NFKD), as specified by [Unicode15]. `<normalized predicate>` may be used to verify the normalization form to which a particular UCS string conforms. Applications may also use `<normalize function>` to enforce a particular `<normal form>`. With the exception of `<normalize function>` and `<normalized predicate>`, the result of any operation on an unnormalized UCS string is implementation-defined.

Conversion of UCS strings from one character set to another is automatic.

Detection of a noncharacter in a UCS-string causes an exception condition to be raised. The detection of an unassigned code point does not.

4.3 Binary strings

4.3.1 Introduction to binary strings

A binary string is a sequence of octets that does not have either a character set or collation associated with it.

A binary string type is described by a binary string type descriptor. A binary string type descriptor contains:

- The name of the data type (BINARY, BINARY VARYING, or BINARY LARGE OBJECT).
- The length or maximum length in octets of the binary string type.

A binary string is assignable only to sites of binary string type. If a store assignment would result in the loss of non-zero octets due to truncation, then an exception condition is raised. If a retrieval assignment would result in the loss of octets due to truncation, then a warning condition is raised.

4.3.2 Binary string comparison

All binary string values are comparable. When binary large object string values are compared, they shall have exactly the same length (in octets) to be considered equal. Binary large object string values can be compared only for equality.

For binary string values other than binary large object string values, it is implementation-defined whether trailing X'00's are considered significant when comparing two binary string values that are otherwise equivalent.

4.3.3 Operations involving binary strings

4.3.3.1 Operators that operate on binary strings and return binary strings

<binary concatenation> is an operator, `||`, that returns a binary string by joining its binary string operands in the order given.

<binary substring function> is a triadic function identical in syntax and semantics to <character substring function> except that its arguments and the returned value are all binary strings.

<binary overlay function> is a function identical in syntax and semantics to <character overlay function> except that the first argument, second argument, and returned value are all binary strings.

<binary trim function> is a function identical in syntax and semantics to <trim function> except that its arguments and the returned value are all binary strings.

4.3.3.2 Other operators involving binary strings

<length expression> returns the length of a given binary string, as an exact numeric value, in octets.

<position expression> when applied to binary strings is identical in syntax and semantics to the corresponding operation on character strings except that the operands are binary strings.

<like predicate> when applied to binary strings is identical in syntax and semantics to the corresponding operation on character strings except that the operands are binary strings.

Binary large object strings cannot be used in:

- Predicates other than <comparison predicate> with an <equals operator> or a <not equals operator>, <quantified comparison predicate> with an <equals operator> or a <not equals operator>, and <exists predicate>.
- <general set function>.
- <group by clause>.
- <order by clause>.
- <unique constraint definition>.
- <referential constraint definition>.
- <select list> of a <query specification> that has a <set quantifier> of DISTINCT.
- UNION, INTERSECT, and EXCEPT.
- Columns used for matching when forming a <joined table>.

4.4 Numbers

4.4.1 Introduction to numbers

A number is either an exact numeric value or an approximate numeric value. Any two numbers are comparable.

A numeric type is described by a numeric type descriptor. A numeric type descriptor contains:

- The name of the specific numeric type (NUMERIC, DECIMAL, SMALLINT, INTEGER, BIGINT, FLOAT, REAL, or DOUBLE PRECISION).
- The implemented precision of the numeric type.
- For an exact numeric type, the implemented scale of the numeric type.
- An indication of whether the precision (and scale) are expressed in decimal or binary terms.
- The name by which the numeric type was declared (but the corresponding full name if INT or DEC was specified).
- The explicit declared precision, if any, of the numeric type.
- For an exact numeric type, the explicit declared scale, if any, of the numeric type.

If <precision> or <scale> is not specified explicitly, then the corresponding element of the descriptor effectively contains the null value.

An SQL-implementation is permitted to regard certain <exact numeric type>s as equivalent, if they have the same precision, scale, and radix, as permitted by the Syntax Rules of [Subclause 6.1, “<data type>”](#). When two or more <exact numeric type>s are equivalent, the SQL-implementation chooses one of these equivalent <exact numeric type>s as the *normal form* representing that equivalence class of <exact numeric type>s. The normal form determines the name of the exact numeric type in the numeric type descriptor.

Similarly, an SQL-implementation is permitted to regard certain <approximate numeric type>s as equivalent, as permitted the Syntax Rules of [Subclause 6.1, “<data type>”](#), in which case the SQL-implementation chooses a *normal form* to represent each equivalence class of <approximate numeric type> and the normal form determines the name of the approximate numeric type.

For every numeric type, the least value is less than zero and the greatest value is greater than zero.

4.4.2 Characteristics of numbers

An exact numeric type has a precision P and a scale S . P is a positive integer that determines the number of significant digits in a particular radix R , where R is either 2 or 10. S is a non-negative integer. Every value of an exact numeric type of scale S is of the form $n \times 10^{-S}$, where n is an integer such that $-R^P \leq n < R^P$.

NOTE 16 — Not every value in that range is necessarily a value of the type in question.

An approximate numeric value consists of a mantissa and an exponent. The mantissa is a signed numeric value, and the exponent is a signed integer that specifies the magnitude of the mantissa. An approximate numeric value has a precision. The precision is a positive integer that specifies the number of significant binary digits

in the mantissa. The value of an approximate numeric value is the mantissa multiplied by a factor determined by the exponent.

An <approximate numeric literal> *ANL* consists of an <exact numeric literal> (called the <mantissa>), the letter 'E' or 'e', and a <signed integer> (called the <exponent>). If M is the value of the <mantissa> and E is the value of the <exponent>, then $M * 10^E$ is the *apparent value* of *ANL*. The actual value of *ANL* is approximately the apparent value of *ANL*, according to implementation-defined rules.

A number is assignable only to sites of numeric type. If an assignment of some number would result in a loss of its most significant digit, an exception condition is raised. If least significant digits are lost, implementation-defined rounding or truncating occurs, with no exception condition being raised. The rules for arithmetic are specified in [Subclause 6.26](#), “<numeric value expression>”.

Whenever an exact or approximate numeric value is assigned to an exact numeric value site, an approximation of its value that preserves leading significant digits after rounding or truncating is represented in the declared type of the target. The value is converted to have the precision and scale of the target. The choice of whether to truncate or round is implementation-defined.

An approximation obtained by truncation of a numeric value N for an <exact numeric type> T is a value V in T such that N is not closer to zero than is V and there is no value in T between V and N .

An approximation obtained by rounding of a numeric value N for an <exact numeric type> T is a value V in T such that the absolute value of the difference between N and the numeric value of V is not greater than half the absolute value of the difference between two successive numeric values in T . If there is more than one such value V , then it is implementation-defined which one is taken.

All numeric values between the smallest and the largest value, inclusive, in a given exact numeric type have an approximation obtained by rounding or truncation for that type; it is implementation-defined which other numeric values have such approximations.

An approximation obtained by truncation or rounding of a numeric value N for an <approximate numeric type> T is a value V in T such that there is no numeric value in T and distinct from that of V that lies between the numeric value of V and N , inclusive.

If there is more than one such value V then it is implementation-defined which one is taken. It is implementation-defined which numeric values have approximations obtained by rounding or truncation for a given approximate numeric type.

Whenever an exact or approximate numeric value is assigned to an approximate numeric value site, an approximation of its value is represented in the declared type of the target. The value is converted to have the precision of the target.

Operations on numbers are performed according to the normal rules of arithmetic, within implementation-defined limits, except as provided for in [Subclause 6.26](#), “<numeric value expression>”.

4.4.3 Operations involving numbers

As well as the usual arithmetic operators, plus, minus, times, divide, unary plus, and unary minus, there are the following functions that return numbers:


- <position expression> (see [Subclause 4.2.3](#), “Operations involving character strings”, and [Subclause 4.3.3](#), “Operations involving binary strings”) takes two strings as arguments and returns an integer.

4.4 Numbers

- <length expression> (see Subclause 4.2.3, “Operations involving character strings”, and Subclause 4.3.3, “Operations involving binary strings”) operates on a string argument and returns an integer.
- <extract expression> (see Subclause 4.6.4, “Operations involving datetimes and intervals”) operates on a datetime or interval argument and returns an exact numeric.
- <cardinality expression> (see Subclause 4.10.5, “Operations involving arrays”, and Subclause 4.10.6, “Operations involving multisets”) operates on a collection argument and returns an integer.
- <absolute value expression> operates on a numeric argument and returns its absolute value in the same most specific type.
- <modulus expression> operates on two exact numeric arguments with scale 0 (zero) and returns the modulus (remainder) of the first argument divided by the second argument as an exact numeric with scale 0 (zero).
- <natural logarithm> computes the natural logarithm of its argument.
- <exponential function> computes the exponential function, that is, e , (the base of natural logarithms) raised to the power equal to its argument.
- <power function> raises its first argument to the power of its second argument.
- <square root> computes the square root of its argument.
- <floor function> computes the greatest integer less than or equal to its argument.
- <ceiling function> computes the least integer greater than or equal to its argument.
- <width bucket function> is a function of four arguments, returning an integer between 0 (zero) and the value of the final argument plus 1 (one), by assigning the first argument to an equi-width partitioning of the range of numbers between the second and third arguments. Values outside the range between the second and third arguments are assigned to either 0 (zero) or the value of the final argument plus 1 (one).

NOTE 17 — The semantics of <width bucket function> are illustrated in Figure 2, “Illustration of WIDTH_BUCKET Semantics”.

Figure 2 — Illustration of WIDTH_BUCKET Semantics



4.5 Boolean types

4.5.1 Introduction to Boolean types

The data type boolean comprises the distinct truth values *True* and *False*. Unless prohibited by a NOT NULL constraint, the boolean data type also supports the truth value *Unknown* as the null value. This specification does not make a distinction between the null value of the boolean data type and the truth value *Unknown* that is the result of an SQL <predicate>, <search condition>, or <boolean value expression>; they may be used interchangeably to mean exactly the same thing.

The boolean data type is described by the boolean data type descriptor. The boolean data type descriptor contains:

— The name of the boolean data type (BOOLEAN).

4.5.2 Comparison and assignment of booleans

All boolean values and SQL truth values are comparable and all are assignable to a site of type boolean. The value *True* is greater than the value *False*, and any comparison involving the null value or an *Unknown* truth value will return an *Unknown* result. The values *True* and *False* may be assigned to any site having a boolean data type; assignment of *Unknown*, or the null value, is subject to the nullability characteristic of the target.

4.5.3 Operations involving booleans

4.5.3.1 Operations on booleans that return booleans

The monadic boolean operator NOT and the dyadic boolean operators AND and OR take boolean operands and produce a boolean result (see Table 11, “Truth table for the AND boolean operator”, and Table 12, “Truth table for the OR boolean operator”).

4.5.3.2 Other operators involving booleans

Every SQL <predicate>, <search condition>, and <boolean value expression> may be considered as an operator that returns a boolean result.

4.6 Datetimes and intervals

4.6.1 Introduction to datetimes and intervals

A datetime data type is described by a datetime data type descriptor. An interval data type is described by an interval data type descriptor.

A datetime data type descriptor contains:

- The name of the specific datetime data type (DATE, TIME WITHOUT TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, TIME WITH TIME ZONE, or TIMESTAMP WITH TIME ZONE).
- The value of the <time fractional seconds precision>, if it is a TIME WITHOUT TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, TIME WITH TIME ZONE, or TIMESTAMP WITH TIME ZONE type.

An interval data type descriptor contains:

- The name of the interval data type (INTERVAL).
- An indication of whether the interval data type is a year-month interval or a day-time interval.
- The <interval qualifier> that describes the precision of the interval data type.

Values of interval data types are always signed.

Every datetime or interval data type has an implied *length in positions*. Let D denote a value in some datetime or interval data type DT . The length in positions of DT is constant for all D . The length in positions is the number of characters from the character set SQL_TEXT that it would take to represent any value in a given datetime or interval data type.

An approximation obtained by rounding of a datetime or interval value D for a datetime type or interval type T is a value V in T such that the absolute value of the difference between D and the numeric value of V is not greater than half the absolute value of the difference between two successive datetime or interval values in T . If there is more than one such value V , then it is implementation-defined which one is taken.

NOTE 18 — No host language for which a binding is specified in this part of ISO/IEC 9075 has a data type corresponding to the datetime and interval data types specified by this part of ISO/IEC 9075. Whenever an application written in one of those languages must transfer values of those data types from the SQL-implementation to the host language program, those values must be transformed to values of some data type for which the host language has a corresponding data type. Typically, the rules of <cast specification> are used to transform values of those data types to values of some character string type. Similar solutions are required when transferring such values from a host program into the SQL-implementation.

4.6.2 Datetimes

Table 2, “Fields in datetime values”, specifies the fields that can make up a datetime value; a datetime value is made up of a subset of those fields. Not all of the fields shown are required to be in the subset, but every field that appears in the table between the first included <primary datetime field> and the last included <primary datetime field> shall also be included. If either <time zone field> is in the subset, then both of them shall be included.

Table 2 — Fields in datetime values

Keyword	Meaning
YEAR	Year
MONTH	Month within year
DAY	Day within month
HOURL	Hour within day
MINUTE	Minute within hour
SECOND	Second and possibly fraction of a second within minute
TIMEZONE_HOUR	Hour value of time zone displacement
TIMEZONE_MINUTE	Minute value of time zone displacement

There is an ordering of the significance of <primary datetime field>s. This is, from most significant to least significant: YEAR, MONTH, DAY, HOUR, MINUTE, and SECOND.

The <primary datetime field>s other than SECOND contain non-negative integer values, constrained by the natural rules for dates using the Gregorian calendar. SECOND, however, can be defined to have a <time fractional seconds precision> that indicates the number of decimal digits maintained following the decimal point in the seconds value, a non-negative exact numeric value.

There are three classes of datetime data types defined within this part of ISO/IEC 9075:

- DATE — contains the <primary datetime field>s YEAR, MONTH, and DAY.
- TIME — contains the <primary datetime field>s HOUR, MINUTE, and SECOND.
- TIMESTAMP — contains the <primary datetime field>s YEAR, MONTH, DAY, HOUR, MINUTE, and SECOND.

Items of type datetime are comparable only if they have the same <primary datetime field>s.

A datetime data type that specifies WITH TIME ZONE is a data type that is *datetime with time zone*, while a datetime data type that specifies WITHOUT TIME ZONE is a data type that is *datetime without time zone*.

The surface of the earth is divided into zones, called time zones, in which every correct clock tells the same time, known as *local time*. Local time is equal to UTC (Coordinated Universal Time) plus the *time zone displacement*, which is an interval value that ranges between INTERVAL '-14:00' HOUR TO MINUTE and INTERVAL '+14:00' HOUR TO MINUTE. The time zone displacement is constant throughout a time zone, changing at the beginning and end of Summer Time, where applicable.

A datetime value, of data type TIME WITHOUT TIME ZONE or TIMESTAMP WITHOUT TIME ZONE, may represent a local time, whereas a datetime value of data type TIME WITH TIME ZONE or TIMESTAMP WITH TIME ZONE represents UTC.

4.6 Datetimes and intervals

On occasion, UTC is adjusted by the omission of a second or the insertion of a “leap second” in order to maintain synchronization with sidereal time. This implies that sometimes, but very rarely, a particular minute will contain exactly 59, 61, or 62 seconds. Whether an SQL-implementation supports leap seconds, and the consequences of such support for date and interval arithmetic, is implementation-defined.

For the convenience of users, whenever a datetime value with time zone is to be implicitly derived from one without (for example, in a simple assignment operation), SQL assumes the value without time zone to be local, subtracts the current default time zone displacement of the SQL-session from it to give UTC, and associates that time zone displacement with the result.

Conversely, whenever a datetime value without time zone is to be implicitly derived from one with, SQL assumes the value with time zone to be UTC, adds the time zone displacement to it to give local time, and the result, without any time zone displacement, is local.

The preceding principles, as implemented by <cast specification>, result in data type conversions between the various datetime data types, as summarized in Table 3, “Datetime data type conversions”.

Table 3 — Datetime data type conversions

	to DATE	to TIME WITHOUT ZONE	to TIME WITH TIME ZONE	to TIMESTAMP WITHOUT TIME ZONE	to TIMESTAMP WITH TIME ZONE
from DATE	<i>trivial</i>	<i>not supported</i>	<i>not supported</i>	Copy year, month, and day; set hour, minute, and second to 0 (zero)	$SV \Rightarrow TS_{w/oTZ}$ $\Rightarrow TS_{w/TZ}$
from TIME WITHOUT TIME ZONE	<i>not supported</i>	<i>trivial</i>	$TV.UTC = SV - STZD$ (modulo 24); $TV.TZ = STZD$	Copy date fields from CURRENT_DATE and time fields from SV	$SV \Rightarrow TS_{w/oTZ}$ $\Rightarrow TS_{w/TZ}$
from TIME WITH TIME ZONE	<i>not supported</i>	$SV.UTC + SV.TZ$ (modulo 24)	<i>trivial</i>	$SV \Rightarrow TS_{w/TZ}$ $\Rightarrow TS_{w/oTZ}$	Copy date fields from CURRENT_DATE and time and time zone fields from SV
from TIMES- TAMP WITHOUT TIME ZONE	Copy date fields from SV	Copy time fields from SV	$SV \Rightarrow TS_{w/TZ}$ $\Rightarrow TIME_{w/TZ}$	<i>trivial</i>	$TV.UTC = SV - STZD$; $TV.TZ = STZD$

	to DATE	to TIME WITHOUT ZONE	to TIME WITH TIME ZONE	to TIMESTAMP WITHOUT TIME ZONE	to TIMESTAMP WITH TIME ZONE
from TIMES- TAMP WITH TIME ZONE	$SV \Rightarrow \text{TSw/oTZ} \Rightarrow \text{DATE}$	$SV \Rightarrow \text{TSw/oTZ} \Rightarrow \text{TIMEw/oTZ}$	Copy time and time zone fields from SV	$SV.UTC + SV.TZ$	<i>trivial</i>
[†] Where <i>SV</i> is the source value, <i>TV</i> is the target value, <i>UTC</i> is the UTC component of <i>SV</i> or <i>TV</i> (if and only if the source or target has time zone), <i>TZ</i> is the timezone displacement of <i>SV</i> or <i>TV</i> (if and only if the source or target has time zone), <i>STZD</i> is the SQL-session default time zone displacement, \Rightarrow means to cast from the type preceding the arrow to the type following the arrow, “TIMEw/TZ” is “TIME WITH TIME ZONE”, “TIMEw/oTZ” is “TIME WITHOUT TIME ZONE”, “TSw/TZ” is “TIMESTAMP WITH TIME ZONE”, and “TSw/oTZ” is “TIMESTAMP WITHOUT TIME ZONE”.					

A datetime is assignable to a site only if the source and target of the assignment are both of type DATE, or both of type TIME (regardless whether WITH TIME ZONE or WITHOUT TIME ZONE is specified or implicit), or both of type TIMESTAMP (regardless whether WITH TIME ZONE or WITHOUT TIME ZONE is specified or implicit).

4.6.3 Intervals

There are two classes of intervals. One class, called *year-month* intervals, has an express or implied datetime precision that includes no fields other than YEAR and MONTH, though not both are required. The other class, called *day-time* intervals, has an express or implied interval precision that can include any fields other than YEAR or MONTH.

Table 4, “Fields in year-month INTERVAL values”, specifies the fields that make up a year-month interval. A year-month interval is made up of a contiguous subset of those fields.

Table 4 — Fields in year-month INTERVAL values

Keyword	Meaning
YEAR	Years
MONTH	Months

Table 5, “Fields in day-time INTERVAL values”, specifies the fields that make up a day-time interval. A day-time interval is made up of a contiguous subset of those fields.

Table 5 — Fields in day-time INTERVAL values

Keyword	Meaning
DAY	Days
HOUR	Hours
MINUTE	Minutes
SECOND	Seconds and possibly fractions of a second

The actual subset of fields that comprise a value of either type of interval is defined by an <interval qualifier> and this subset is known as the precision of the value.

Within a value of type interval, the first field is constrained only by the <interval leading field precision> of the associated <interval qualifier>. [Table 6, “Valid values for fields in INTERVAL values”](#), specifies the constraints on subsequent field values.

Table 6 — Valid values for fields in INTERVAL values

Keyword	Valid values of INTERVAL fields
YEAR	Unconstrained except by <interval leading field precision>
MONTH	Months (within years) (0-11)
DAY	Unconstrained except by <interval leading field precision>
HOUR	Hours (within days) (0-23)
MINUTE	Minutes (within hours) (0-59)
SECOND	Seconds (within minutes) (0-59.999...)

Values in interval fields other than SECOND are integers and have precision 2 when not the first field. SECOND, however, can be defined to have an <interval fractional seconds precision> that indicates the number of decimal digits maintained following the decimal point in the seconds value. When not the first field, SECOND has a precision of 2 places before the decimal point.

Fields comprising an item of type interval are also constrained by the definition of the Gregorian calendar.

Year-month intervals are comparable only with other year-month intervals. If two year-month intervals have different interval precisions, they are, for the purpose of any operations between them, effectively converted to the same precision by appending new <primary datetime field>s to either the most significant end of one interval, the least significant end of one interval, or both. New least significant <primary datetime field>s are assigned a value of 0 (zero). When it is necessary to add new most significant datetime fields, the associated value is effectively converted to the new precision in a manner obeying the natural rules for dates and times associated with the Gregorian calendar.

Day-time intervals are comparable only with other day-time intervals. If two day-time intervals have different interval precisions, they are, for the purpose of any operations between them, effectively converted to the same precision by appending new <primary datetime field>s to either the most significant end of one interval or the least significant end of one interval, or both. New least significant <primary datetime field>s are assigned a value of 0 (zero). When it is necessary to add new most significant datetime fields, the associated value is effectively converted to the new precision in a manner obeying the natural rules for dates and times associated with the Gregorian calendar.

4.6.4 Operations involving datetimes and intervals

Table 7, “Valid operators involving datetimes and intervals”, specifies the declared types of arithmetic expressions involving datetime and interval operands.

Table 7 — Valid operators involving datetimes and intervals

Operand 1	Operator	Operand 2	Result Type
Datetime	–	Datetime	Interval
Datetime	+ or –	Interval	Datetime
Interval	+	Datetime	Datetime
Interval	+ or –	Interval	Interval
Interval	* or /	Numeric	Interval
Numeric	*	Interval	Interval

Arithmetic operations involving values of type datetime or interval obey the natural rules associated with dates and times and yield valid datetime or interval results according to the Gregorian calendar.

Operations involving values of type datetime require that the datetime values be comparable. Operations involving values of type interval require that the interval values be comparable.

Operations involving a datetime and an interval preserve the time zone of the datetime operand. If the datetime operand does not include a time zone displacement, then the result has no time zone displacement.

<overlaps predicate> uses the operator OVERLAPS to determine whether or not two chronological periods overlap in time. A chronological period is specified either as a pair of datetimes (starting and ending) or as a starting datetime and an interval. If the length of the period is greater than 0 (zero), then the period consists of all points of time greater than or equal to the lower endpoint, and less than the upper endpoint. If the length of the period is equal to 0 (zero), then the period consists of a single point in time, the lower endpoint. Two periods overlap if they have at least one point in common.

<extract expression> operates on a datetime or interval and returns an exact numeric value representing the value of one component of the datetime or interval.

<interval absolute value function> operates on an interval argument and returns its absolute value in the same most specific type.

4.7 User-defined types

This Subclause is modified by Subclause 4.8, “User-defined types”, in ISO/IEC 9075-13.

4.7.1 Introduction to user-defined types

This Subclause is modified by Subclause 4.8.1, “Introduction to user-defined types”, in ISO/IEC 9075-13.

13 A user-defined type is a schema object, identified by a <user-defined type name>. The definition of a user-defined type specifies a representation for values of that type. In some cases, known as *distinct types*, the representation is a single predefined type, known as the *source type*; in others, *structured types*, it consists of a list of attribute definitions. Although the attribute definitions are said to define the representation of the user-defined type, in fact they implicitly define certain functions (observer functions and mutator functions) that are part of the interface of the user-defined type; physical representations of user-defined type values are implementation-dependent.

The definition of a user-defined type may include a <method specification list> consisting of one or more <method specification>s. A <method specification> is either an <original method specification> or an <overriding method specification> (in which case the type being defined must be a structured type). Each <original method specification> specifies:

- The <method name>.
- The <specific method name>.
- The <SQL parameter declaration list>.
- The <returns data type>.
- The <result cast from type> (if any).
- Whether the method is type-preserving.
- The <language clause>.
- If the language is not SQL, then the <parameter style>.
- Whether STATIC or CONSTRUCTOR is specified.
- Whether the method is deterministic.
- Whether the method possibly modifies SQL-data, possibly reads SQL-data, possibly contains SQL, or does not possibly contain SQL.
- Whether the method should be evaluated as the null value whenever any argument is the null value, without actually invoking the method.
- The creation timestamp.

Each <overriding method specification> specifies the <method name>, the <specific method name>, the <SQL parameter declaration list> and the <returns data type>. For each <overriding method specification>, there shall be an <original method specification> with the same <method name> and <SQL parameter declaration list> in some proper supertype of the user-defined type. Every SQL-invoked method in a schema shall correspond

to exactly one <original method specification> or <overriding method specification> associated with some user-defined type existing in that schema.

A method *M* that corresponds to an <original method specification> in the definition of a structured type *T1* is an *original method* of *T1*. A method *M* that corresponds to an <overriding method specification> in the definition of *T1* is an *overriding method* of *T1*.

A method *M* is a *method of type T1* if one of the following holds:

- *M* is an original method of *T1*.
- *M* is an overriding method of *T1*.
- There is a proper supertype *T2* of *T1* such that *M* is an original or overriding method of *T2* and such that there is no method *M3* such that *M3* has the same <method name> and <SQL parameter declaration list> as *M* and *M3* is an original method or overriding method of a type *T3* such that *T2* is a proper supertype of *T3* and *T3* is a supertype of *T1*.

If *T1* is a subtype of *T2* and *M1* is a method of *T1* such that there exists a method *M2* of *T2* such that *M1* and *M2* have the same <method name> and the same unaugmented <SQL parameter declaration list>, then *M1* is an *inherited method* of *T1* from *T2*.

4.7.2 User-defined type descriptor

This Subclause is modified by Subclause 4.8.2, “User-defined type descriptor”, in ISO/IEC 9075-13.

13 A user-defined type is described by a user-defined type descriptor. A user-defined type descriptor contains:

- The name of the user-defined type (<user-defined type name>). This is the type designator of that type, used in type precedence lists (see Subclause 9.5, “Type precedence list determination”).
- An indication of whether the user-defined type is a structured type or a distinct type.
- The ordering form for the user-defined type (EQUALS, FULL, or NONE).
- **13** The ordering category for the user-defined type (RELATIVE, MAP, or STATE).
- A <specific routine designator> identifying the ordering function, depending on the ordering category.
- If the user-defined type is a direct subtype of another user-defined type, then the name of that user-defined type.
- If the representation is a predefined data type, then the descriptor of that type; otherwise the attribute descriptor of every originally-defined attribute and every inherited attribute of the user-defined type.
- An indication of whether the user-defined type is instantiable or not instantiable.
- An indication of whether the user-defined type is final or not final.
- **13** The transform descriptor of the user-defined type.
- If the user-defined type is a structured type, then:
 - Whether the referencing type of the structured type has a user-defined representation, a derived representation, or a system-defined representation.

4.7 User-defined types

- If user-defined representation is specified, then the type descriptor of the representation type of the referencing type of the structured type; otherwise, if derived representation is specified, then the list of attributes.

NOTE 19 — “user-defined representation”, “derived representation”, and “system-defined representation” of a reference type are defined in Subclause 4.9, “Reference types”.

— If the <method specification list> is specified, then for each <method specification> contained in <method specification list>, a *method specification descriptor* that includes:

- The <method name>.
- The <specific method name>.
- The <SQL parameter declaration list> augmented to include the implicit first parameter with parameter name SELF.
- The <language name>.
- If the <language name> is not SQL, then the <parameter style>.
- The <returns data type>.
- The <result cast from type>, if any.
- An indication as to whether the <method specification> is an <original method specification> or an <overriding method specification>.
- If the <method specification> is an <original method specification>, then an indication of whether STATIC or CONSTRUCTOR is specified.
- An indication whether the method is deterministic.
- An indication whether the method possibly modifies SQL-data, possibly reads SQL-data, possibly contains SQL, or does not possibly contain SQL.
- An indication whether the method should not be invoked if any argument is the null value, in which case the value of the method is the null value.

NOTE 20 — The characteristics of an <overriding method specification> other than the <method name>, <SQL parameter declaration list>, and <returns data type> are the same as the characteristics for the corresponding <original method specification>.

4.7.3 Observer functions and mutator functions

Corresponding to every attribute of every structured type is exactly one implicitly-defined observer function and exactly one implicitly-defined mutator function. These are both SQL-invoked functions. Further, the mutator function is a type-preserving function.

Let A be the name of an attribute of structured type T and let AT be the data type of A . The signature of the observer function for this attribute is FUNCTION $A(T)$ and its result data type is AT . The signature of the mutator function for this attribute is FUNCTION $A(T \text{ RESULT}, AT)$ and its result data type is T .

Let V be a value in data type T and let AV be a value in data type AT . The invocation $A(V, AV)$ returns MV such that “ $A(MV)$ is identical to AV ” and for every attribute A' ($A' \neq A$) of T , “ $A'(MV)$ is identical to $A'(V)$ ”. The most specific type of MV is the most specific type of V .

4.7.4 Constructors

Associated with each structured type ST is one implicitly defined *constructor function*, if and only if ST is instantiable.

Let TN be the name of a structured type T . The signature of the constructor function for T is $TN()$ and its result data type is T . The invocation $TN()$ returns a value V such that V is not the null value and, for every attribute A of T , $A(V)$ returns the default value of A . The most specific type of V is T .

For every structured type ST that is instantiable, zero or more SQL-invoked constructor methods can be specified. The names of those methods shall be equivalent to the name of the type for which they are specified.

NOTE 21 — SQL-invoked constructor methods are original methods that cannot be overloaded. An SQL-invoked constructor method and a regular SQL-invoked function may exist such that they have equivalent routine names, the types of the first parameter of the method's augmented parameter list and the function's parameter list are the same, and the types of the corresponding remaining parameters (if any) are identical according to the Syntax Rules of Subclause 9.18, "Data type identity".

4.7.5 Subtypes and supertypes

As a consequence of the <subtype clause> of <user-defined type definition>, two structured types T_a and T_b that are not compatible can be such that T_a is a subtype of T_b . See Subclause 11.42, "<user-defined type definition>".

A type T_a is a *direct subtype* of a type T_b if T_a is a proper subtype of T_b and there does not exist a type T_c such that T_c is a proper subtype of T_b and a proper supertype of T_a .

A type T_a is a subtype of type T_b if one of the following pertains:

- T_a and T_b are compatible;
- T_a is a direct subtype of T_b ; or
- T_a is a subtype of some type T_c and T_c is a direct subtype of T_b .

By the same token, T_b is a supertype of T_a and is a direct supertype of T_a in the particular case where T_a is a direct subtype of T_b .

If T_a is a subtype of T_b and T_a and T_b are not compatible, then T_a is a proper subtype of T_b and T_b is a proper supertype of T_a . A type cannot be a proper supertype of itself.

A type with no proper supertypes is a maximal supertype. A type with no proper subtypes is a leaf type.

Let T_a be a maximal supertype and let T be a subtype of T_a . The set of all subtypes of T_a (which includes T_a itself) is called a *subtype family* of T or (equivalently) of T_a . A subtype family is not permitted to have more than one maximal supertype.

Every value in a type T is a value in the direct supertype, if any, of T . A value V in type T has exactly one most specific type MST such that MST is a subtype of T and V is not a value in any proper subtype of MST . The most specific type of value need not be a leaf type. For example, a type structure might consist of a type PERSON that has STUDENT and EMPLOYEE as its two subtypes, while STUDENT has two direct subtypes UG_STUDENT and PG_STUDENT. The invocation STUDENT() of the constructor function for STUDENT returns a value whose most specific type is STUDENT, which is not a leaf type.

4.7 User-defined types

If T_a is a subtype of T_b , then a value in T_a can be used wherever a value in T_b is expected. In particular, a value in T_a can be stored in a column of type T_b , can be substituted as an argument for an input SQL parameter of data type T_b , and can be the value of an invocation of an SQL-invoked function whose result data type is T_b .

A type T is said to be the *minimal common supertype* of a set of types S if T is a supertype of every type in S and a subtype of every type that is a supertype of every type in S .

NOTE 22 — Because a subtype family has exactly one maximal supertype, if two types have a common subtype, they shall also have a minimal common supertype. Thus, for every set of types drawn from the same subtype family, there is some member of that family that is the minimal common supertype of all of the types in that set.

If a structured type ST is defined to be not instantiable, then the most specific type of every value in ST is necessarily of some proper subtype of ST .

If a user-defined type UDT is defined to be final, then UDT has no proper subtypes. As a consequence, the most specific type of every value in UDT is necessarily UDT .

Users shall have the UNDER privilege on a type before they can define any direct subtypes of it. A type can have more than one direct subtype. A user-defined type or a reference type can have at most one direct supertype. A row type can have more than one direct supertype.

A <user-defined type definition> for type T can include references to components of the direct supertype, if any, of T . Effectively, components of the representation of the direct supertype are copied to the subtype's representation.

4.7.6 User-defined type comparison and assignment

This Subclause is modified by Subclause 4.8.3, “User-defined type comparison and assignment”, in ISO/IEC 9075-13.

Let *comparison type* of a user-defined type T_a be the user-defined type T_b that satisfies all the following conditions:

- 1) The type designator of T_b is in the type precedence list of T_a .
- 2) The user-defined type descriptor of T_b includes an ordering form that is EQUALS or FULL.
- 3) The descriptor of no type T_c whose type designator precedes that of T_b in the type precedence list of T_a includes an ordering form that includes EQUALS or FULL.

If there is no such type T_b , then T_a has no comparison type.

Let *comparison form* of a user-defined type T_a be the ordering form included in the user-defined type descriptor of the comparison type of T_a .

Let *comparison category* of a user-defined type T_a be the ordering category included in the user-defined type descriptor of the comparison type of T_a .

13 Let *comparison function* of a user-defined type T_a be the ordering function included in the user-defined type descriptor of the comparison type of T_a .

13 Two values $V1$ and $V2$ whose most specific types are user-defined types $T1$ and $T2$ are comparable if and only if $T1$ and $T2$ are in the same subtype family and each have some comparison type $CT1$ and $CT2$, respectively. $CT1$ and $CT2$ constrain the comparison forms and comparison categories of $T1$ and $T2$ to be the same and to be the same as those of all their supertypes. If the comparison category is either STATE or RELATIVE, then $T1$ and $T2$ are constrained to have the same comparison function; if the comparison category is MAP, they are not constrained to have the same comparison function.

NOTE 23 — Explicit cast functions or attribute comparisons can be used to make both values of the same subtype family or to perform the comparisons on attributes of the user-defined types.

NOTE 24 — “Subtype” and “subtype family” are defined in Subclause 4.7.5, “Subtypes and supertypes”.

If there is no appropriate user-defined cast function, then an expression E whose declared type is some user-defined type $UDT1$ is assignable to a site S whose declared type is some user-defined type $UDT2$ if and only if $UDT1$ is a subtype of $UDT2$. The effect of the assignment of E to S is that the value of S is V , obtained by the evaluation of E . The most specific type of V is some subtype of $UDT1$, possibly $UDT1$ itself, while the declared type of S remains $UDT2$.

An expression whose declared type is some distinct type whose source type is SDT is assignable to any site whose declared type is SDT because of the implicit cast functions created by the General Rules of Subclause 11.42, “<user-defined type definition>”. Similarly, an expression whose declared type is some pre-defined data type SDT is assignable to any site whose declared type is some distinct type whose source type is SDT .

4.7.7 Transforms for user-defined types

Transforms are SQL-invoked functions that are automatically invoked when values of user-defined types are transferred from SQL-environment to host languages or vice-versa.

A transform is associated with a user-defined type. A transform identifies a list of *transform groups* of up to two SQL-invoked functions, called the *transform functions*, each identified by a group name. The group name of a transform group is an <identifier> such that no two transform groups for a transform have the same group name. The two transform functions are:

- **from-sql function** — This SQL-invoked function maps the user-defined type value into a value of an SQL pre-defined type, and gets invoked whenever a user-defined type value is passed to a host language program or an external routine.
- **to-sql function** — This SQL-invoked function maps a value of an SQL predefined type to a value of a user-defined type and gets invoked whenever a user-defined type value is supplied by a host language program or an external routine.

A transform is defined by a <transform definition>. A transform is described by a *transform descriptor*. A transform descriptor includes a possibly empty list of *transform group descriptors*, where each transform group descriptor includes:

- The group name of the transform group.
- The specific name of the from-sql function, if any, associated with the transform group.
- The specific name of the to-sql function, if any, associated with the transform group.

4.8 Row types

A row type is a sequence of (<field name> <data type>) pairs, called *fields*. It is described by a row type descriptor. A row type descriptor consists of the field descriptor of every field of the row type.

The most specific type of a row of a table is a row type. In this case, each column of the row corresponds to the field of the row type that has the same ordinal position as the column.

Row type $RT2$ is a subtype of data type $RT1$ if and only if $RT1$ and $RT2$ are row types of the same degree and, in every n -th pair of corresponding field definitions, $FD1_n$ in $RT1$ and $FD2_n$ in $RT2$, the <field name>s are equivalent and the <data type> of $FD2_n$ is a subtype of the <data type> of $FD1_n$.

A value of row type $RT1$ is assignable to a site of row type $RT2$ if and only if the degree of $RT1$ is the same as the degree of $RT2$ and every field in $RT1$ is assignable to the field in the same ordinal position in $RT2$.

A value of row type $RT1$ is comparable with a value of row type $RT2$ if and only if the degree of $RT1$ is the same as the degree of $RT2$ and every field in $RT1$ is comparable with the field in the same ordinal position in $RT2$.

4.9 Reference types

4.9.1 Introduction to reference types

A *REF value* is a value that references a row in a *referenceable table* (see Subclause 4.14.5, “Referenceable tables, subtables, and supertables”). A referenceable table is necessarily also a *typed table* (that is, it has an associated structured type from which its row type is derived).

Given a structured type T , the REF values that can reference rows in typed tables defined on T collectively form a certain data type RT known as a *reference type*. RT is the *referencing type* of T and T is the *referenced type* of RT .

Let TN be name of T . The type designator of RT is $REF(TN)$.

Values of two reference types are comparable if the referenced types of their declared types have some common supertype.

An expression E whose declared type is some reference type $RT1$ is assignable to a site S whose declared type is some reference type $RT2$ if and only if the referenced type of $RT1$ is a subtype of the referenced type of $RT2$. The effect of the assignment of E to S is that the value of S is V , obtained by the evaluation of E . The most specific type of V is some subtype of $RT1$, possibly $RT1$ itself, while the declared type of S remains $RT2$.

A site RS that is occupied by a REF value might have a *scope*, which determines the effect of an invocation of <reference resolution> RR on the value at RS . A scope is specified as a table name STN and consists at any time of every row in the table ST identified by STN . ST is the *scoped table* of RR . The scope of RS is specified in the declared type of RS . If no scope is specified in the declared type of RS , then <reference resolution> is not available.

A reference type is described by a reference type descriptor. The reference type descriptor for RT includes:

— The type designator of RT .

— The name of the referenceable table, if any, that is the scope of *RT*.

In a host variable, a REF value is materialized as an *N*-octet value, where *N* is implementation-defined.

Reference type *RT2* is a *subtype* of data type *RT1* (equivalently, *RT1* is a *supertype* of *RT2*) if and only if *RT1* is a reference type and the referenced type of *RT2* is a subtype of the referenced type of *RT1*.

Every value in a reference type *RT* is a value in every supertype of *RT*. A value *V* in type *RT* has exactly one most specific type *MST* such that *MST* is a subtype of *RT* and *V* is not a value in any proper subtype of *MST*.

A reference type has a *user-defined representation* if its referenced type is defined by a <user-defined type definition> that specifies <user-defined representation>. A reference type has a *derived representation* if its referenced type is defined by a <user-defined type definition> that specifies <derived representation>. A reference type has a *system-defined representation* if it does not have a user-defined representation or a derived representation.

4.9.2 Operations involving references

An operation is provided that takes a REF value and returns the value that is held in a column of the site identified by the REF value (see Subclause 6.20, “<dereference operation>”). If the REF value identifies no site, perhaps because a site it once identified has been destroyed, then the null value is returned.

An operation is provided that takes a REF value and returns a value of the referenced type; that value is constructed from the values of the columns of the site identified by that REF value (see Subclause 6.22, “<reference resolution>”). An operation is also provided that takes a REF value and returns a value acquired from invoking an SQL-invoked method on a value of the referenced type; that value is constructed from the values of the columns of the site identified by that REF value (see Subclause 6.21, “<method reference>”).

4.10 Collection types

4.10.1 Introduction to collection types

A *collection* is a composite value comprising zero or more *elements*, each a value of some data type *DT*. If the elements of some collection *C* are values of *DT*, then *C* is said to be a collection of *DT*. The number of elements in *C* is the *cardinality* of *C*. The term “element” is not further defined in this part of ISO/IEC 9075. The term “collection” is generic, encompassing various kinds of collection in connection with each of which, individually, this part of ISO/IEC 9075 defines primitive type constructors and operators. This part of ISO/IEC 9075 supports two kinds of collection types, arrays and multisets.

A specific <collection type> *CT* is a <data type> specified by pairing a keyword *KC* (either ARRAY or MULTiset) with a specific data type *EDT*. In addition, a maximum cardinality may optionally be specified for arrays. Every element of every possible value of *CT* is a value of *EDT* and is permitted to be, more specifically, of some subtype of *EDT*. *EDT* is termed the *element type* of *CT*. *KC* specifies the kind of collection, such as ARRAY or MULTiset, that every value of *CT* is, and thus determines the operators that are available for operating on or returning values of *CT*.

Let *EDTN* be the type designator of *EDT*. The type designator of *CT* is *EDTN KC*.

4.10 Collection types

A *collection type descriptor* describes a collection type. The collection type descriptor for *CT* includes:

- The type designator of *CT*.
- The descriptor of the element type of *CT*.
- An indication of the kind of the collection of *CT*: ARRAY or MULTISSET.
- If *CT* is an array type, the maximum number of elements of *CT*.

Collection type *CT2* is a subtype of data type *CT1* (equivalently, *CT1* is a supertype of *CT2*) if and only if *CT1* is the same kind of collection as *CT2* and the element type of *CT2* is a subtype of the element type of *CT1*.

4.10.2 Arrays

An *array* is a collection *A* in which each element is associated with exactly one ordinal position in *A*. If *n* is the cardinality of *A*, then the ordinal position *p* of an element is an integer in the range $1 \text{ (one)} \leq p \leq n$. If *EDT* is the element type of *A*, then *A* can thus be considered as a function of the integers in the range $1 \text{ (one)} \text{ to } n$ into *EDT*.

An array site *AS* has a maximum cardinality *m*. The cardinality *n* of an array occupying *AS* is constrained not to exceed *m*.

An *array type* is a <collection type>. If *AT* is some array type with element type *EDT*, then every value of *AT* is an array of *EDT*.

Let *A1* and *A2* be arrays of *EDT*. *A1* and *A2* are identical if and only if *A1* and *A2* have the same cardinality *n* and if, for all *i* in the range $1 \text{ (one)} \leq i \leq n$, the element at ordinal position *i* in *A1* is identical to the element at ordinal position *i* in *A2*.

4.10.3 Multisets

A multiset is an unordered collection. Since a multiset is unordered, there is no ordinal position to reference individual elements of a multiset.

A multiset type is a <collection type>. If *MT* is some multiset type with element type *EDT*, then every value of *MT* is a multiset of *EDT*.

Let *M1* and *M2* be multisets of *EDT*. *M1* and *M2* are identical if and only if *M1* and *M2* have the same cardinality *n*, and for each element *x* in *M1*, the number of elements of *M1* that are identical to *x*, including *x* itself, equals the number of elements of *M2* that are identical to *x*.

Let *n1* be the cardinality of *M1* and let *n2* be the cardinality of *M2*. *M1* is a submultiset of *M2* if, for each element *x* of *M1*, the number of elements of *M1* that are not distinct from *x*, including *x* itself, is less than or equal to the number of elements of *M2* that are not distinct from *x*.

4.10.4 Collection comparison and assignment

Two collections are comparable if and only if they are of the same kind of collection (ARRAY or MULTISSET) and their element types are comparable.

A value of collection type $CT1$ is assignable to a site of collection type $CT2$ if and only if $CT1$ is the same kind of collection (ARRAY or MULTISSET) as $CT2$ and the element type of $CT1$ is assignable to the element type of $CT2$.

The array types have a defined *element order*. Comparisons are defined in terms of the element order of the arrays. The element order defines the pairs of corresponding elements from the arrays being compared. The element order of an array is implicitly defined by the ordinal position of its elements.

In the case of comparison of two arrays C and D , the elements are compared pairwise in element order. $C = D$ is True if and only if C and D have the same cardinality and every pair of elements are equal.

Two multisets C and D of comparable element types are equal if they have the same cardinality N and there exist an enumeration CE_j , $1 \text{ (one)} \leq j \leq N$ of the elements of C and an enumeration DE_j , $1 \text{ (one)} \leq j \leq N$ of the elements of D such that for all j , $CE_j = DE_j$.

4.10.5 Operations involving arrays

4.10.5.1 Operators that operate on array values and return array elements

<array element reference> is an operation that returns the array element in the specified position within the array.

4.10.5.2 Operators that operate on array values and return array values

<array concatenation> is an operation that returns the array value made by joining its array value operands in the order given.

4.10.5.3 Operators that operate on array values and return numbers

<cardinality expression> is an operation that returns the cardinality of the array.

4.10.6 Operations involving multisets

4.10.6.1 Operators that operate on multisets and return multiset elements

<multiset element reference> is an operation that returns the value of the element of a multiset, if the multiset has only one element.

4.10.6.2 Operators that operate on multisets and return multisets

<multiset set function> is an operation that returns the multiset obtained by removing duplicates from a multiset.

MULTISET UNION is an operator that computes the union of two multisets. There are two variants, specified using ALL or DISTINCT, to either retain duplicates or remove duplicates.

MULTISET INTERSECT is an operator that computes the intersection of two multisets. There are two variants, ALL and DISTINCT. The variant specified by ALL places in the result as many instances of each value as the minimum number of instances of that value in either operand. The variant specified by DISTINCT removes duplicates from the result.

MULTISET EXCEPT is an operator that computes the multiset difference of two multisets. There are two variants, ALL and DISTINCT. The variant specified by ALL places in the result a number of instances of a value, equal to the number of instances of the value in the first operand minus the number of instances of the value in the second operand. The variant specified by DISTINCT removes duplicates from the result.

4.10.6.3 Operators that operate on multiset values and return numbers

<cardinality expression> is an operation that returns the cardinality of the multiset.

4.11 Data conversions

This Subclause is modified by Subclause 4.3, “Data conversions”, in ISO/IEC 9075-14.

Implicit type conversion can occur in expressions, fetch operations, single row select operations, inserts, deletes, and updates. Explicit type conversions can be specified by the use of the CAST operator.

Explicit data conversions can be specified by a *CAST operator*. A CAST operator defines how values of a source data type are converted into a value of a target data type according to the Syntax Rules and General Rules of Subclause 6.12, “<cast specification>”. Data conversions between predefined data types and between constructed types are defined by the rules of this part of ISO/IEC 9075. Data conversions between a user-defined type and another data type are defined by a user-defined cast.

14 A user-defined cast identifies an SQL-invoked function, called the *cast function*, that has one SQL parameter whose declared type is the same as the source data type and a result data type that is the target data type. A cast function may optionally be specified to be implicitly invoked whenever values are assigned to targets of its result data type. Such a cast function is called an *implicitly invocable* cast function.

A user-defined cast is defined by a <user-defined cast definition>. A user-defined cast has a user-defined cast descriptor that includes:

- The name of the source data type.
- The name of the target data type.
- The specific name of the SQL-invoked function that is the cast function.
- An indication as to whether the cast function is implicitly invocable.

When a value V of declared type TV is assigned to a target T of declared type TT , a user-defined cast function $UDCF$ is said to be an *appropriate user-defined cast function* if and only if all of the following are true:

- The descriptor of $UDCF$ indicates that $UDCF$ is implicitly invocable.
- The type designator of the declared type DTP of the only SQL parameter P of $UDCF$ is in the type precedence list of TV .
- The result data type of $UDCF$ is TT .
- No other user-defined cast function $UDCQ$ with an SQL parameter Q with declared type TQ that precedes DTP in the type precedence list of TV is an appropriate user-defined cast function to assign V to T .

An SQL procedure statement S is said to be *dependent on* an appropriate user-defined cast function $UDCF$ if and only if all of the following are true:

- S is a <select statement: single row>, <insert statement>, <update statement: positioned>, <update statement: searched>, or <merge statement>.
- $UDCF$ is invoked during a store or retrieval assignment operation that is executed during the execution of S and $UDCF$ is not executed during the invocation of an SQL-invoked function that is invoked during the execution of S .

4.12 Domains

A domain is a set of permissible values. A domain is defined in a schema and is identified by a <domain name>. The purpose of a domain is to constrain the set of valid values that can be stored in a column of a base table by various operations.

A domain definition specifies a data type. It may also specify a <domain constraint> that further restricts the valid values of the domain and a <default clause> that specifies the value to be used in the absence of an explicitly specified value or column default.

A domain is described by a domain descriptor. A domain descriptor includes:

- The name of the domain.
- The data type descriptor of the data type of the domain.
- The value of <default option>, if any, of the domain.
- The domain constraint descriptors of the domain constraints, if any, of the domain.

4.13 Columns, fields, and attributes

This Subclause is modified by Subclause 4.9, “Columns, fields, and attributes”, in ISO/IEC 9075-9.

The terms *column*, *field*, and *attribute* refer to structural components of tables, row types, and structured types, respectively, in analogous fashion. As the structure of a table consists of one or more columns, so does the structure of a row type consist of one or more fields and that of a structured type one or more attributes. Every structural element, whether a column, a field, or an attribute, is primarily a name paired with a declared type. The elements of a structure are ordered. Elements in different positions in the same structure can have the same declared type but not the same name. Although the elements of a structure are distinguished from each other by name, in some circumstances the compatibility of two structures (for the purpose at hand) is determined solely by considering the declared types of each pair of elements at the same ordinal position.

A table (see Subclause 4.14, “Tables”) is defined on one or more columns and consists of zero or more rows. A column has a name and a declared type. Each row in a table has exactly one value for each column. Each value in a row is a value in the declared type of the corresponding column. A column is either *updatable* or *not updatable*.

NOTE 25 — The declared type includes the null value and values in proper subtypes of the declared type.

Every column has a *nullability characteristic* that indicates whether the value from that column can be the null value. A nullability characteristic is either *known not nullable* or *possibly nullable*.

Let *C* be a column of a base table *T*. *C* is *known not nullable* if and only if at least one of the following is true:

- There exists at least one constraint *NNC* that is not deferrable and that simply contains a <search condition> that is a <boolean value expression> that is a readily-known-not-null condition for *C*.
- *C* is based on a domain that has a domain constraint that is not deferrable and that simply contains a <search condition> that is a <boolean value expression> that is a readily-known-not-null condition for *VALUE*.
- *C* is a unique column of a nondeferrable unique constraint that is a PRIMARY KEY.
- *C* is the self-referencing column of *T*.
- The SQL-implementation is able to deduce that the <search condition> “*C* IS NULL” can never be *True* when applied to a row in *T* through some additional implementation-defined rule or rules.

NOTE 26 — The notions “known not nullable” and “known-not-null condition” are examples of such additional rules that an SQL-implementation may adopt. If an SQL-implementation does adopt those rules (and possibly others), then the SQL-implementation may claim support for Feature T101, “Enhanced nullability determination”.

The nullability characteristic of a column of a derived table is defined by the the Syntax Rules of Subclause 7.7, “<joined table>”, Subclause 7.12, “<query specification>”, and Subclause 7.13, “<query expression>”.

A column *C* is described by a column descriptor. A column descriptor includes:

- The name of the column.
- Whether the name of the column is an implementation-dependent name.
- If the column is based on a domain, then the name of that domain; otherwise, the data type descriptor of the declared type of *C*.
- The value of <default option>, if any, of *C*.
- The nullability characteristic of *C*.

- The ordinal position of C within the table that contains it.
- An indication of whether C is updatable or not.
- An indication of whether C is a self-referencing column of a base table or not.
- An indication of whether C is an identity column or not.
- If C is an identity column, then an indication of whether values are always generated or generated by default.
- If C is an identity column, then the *start value* of C .
- If C is an identity column, then the descriptor of the internal sequence generator for C .

NOTE 27 — Identity columns and the meaning of “start value” are described in Subclause 4.14.7, “Identity columns”.

- An indication of whether C is a generated column or not.
- If C is a generated column, then the generation expression of C .

NOTE 28 — Generated columns and the meaning of “generation expression” are described in Subclause 4.14.8, “Base columns and generated columns”.

An attribute A is described by an attribute descriptor. An attribute descriptor includes:

- The name of the attribute.
- The data type descriptor of the declared type of A .
- The ordinal position of A within the structured type that contains it.
- The value of the implicit or explicit <attribute default> of A .
- The name of the structured type defined by the <user-defined type definition> that defines A .

A field F is described by a field descriptor. A field descriptor includes:

- The name of the field.
- The data type descriptor of the declared type of F .
- The ordinal position of F within the row type that simply contains it.

4.14 Tables

This Subclause is modified by Subclause 4.2, “Tables”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 4.10, “Tables”, in ISO/IEC 9075-9.

4.14.1 Introduction to tables

A table is a collection of rows having one or more columns. A row is a value of a row type. Every row of the same table has the same row type. The value of the i -th field of every row in a table is the value of the i -th column of that row in the table. The row is the smallest unit of data that can be inserted into a table and deleted from a table.

A table *T2* is *part of* a column *C* of a table *T1* if setting the value of *T1.C* to a null value (ignoring any constraints or triggers defined on *T1* or *T1.C*) would cause *T2* to disappear.

The most specific type of a row is a row type. All rows of a table are of the same row type and this is called the *row type* of that table.

The degree of a table, and the degree of each of its rows, is the number of columns of that table. The number of rows in a table is its cardinality. A table whose cardinality is 0 (zero) is said to be *empty*.

4.14.2 Types of tables

This Subclause is modified by Subclause 4.2.1, “Types of tables”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 4.10.1, “Types of tables”, in ISO/IEC 9075-9.

09 A table is either a base table, a derived table, a transient table, or a viewed table. A base table is either a persistent base table, a global temporary table, a created local temporary table, or a declared local temporary table.

A persistent base table is a named table defined by a <table definition> that does not specify TEMPORARY.

A derived table is a table derived directly or indirectly from one or more other tables by the evaluation of an expression, such as a <query expression> or <table expression>. A <query expression> can contain an optional <order by clause>. The ordering of the rows of the table specified by the <query expression> is guaranteed only for the <query expression> that immediately contains the <order by clause>.

A <query expression> can also optionally contain a <fetch first clause>, which may limit the cardinality of the derived table. If a <query expression> contains both an <order by clause> and a <fetch first clause>, then the rows in the derived table are first sorted according to the <order by clause> and then limited to the number of rows specified in the <fetch first clause>.

NOTE 29 — If the cardinality of the result of an evaluation of a <query expression> is less than the limit specified by a <fetch first clause>, then the <fetch first clause> has no effect on the derived table.

A viewed table is a named derived table defined by a <view definition>. A viewed table is sometimes called a *view*. Base tables and views are identified by <table name>s. The same <table name>, in its fully qualified form, cannot be used for both a base table and a view.

A <query specification>, <table value constructor>, <explicit table>, or <query expression> contained in a <view definition> is called a *view component*.

A transient table is a named table that may come into existence implicitly during the evaluation of a <query expression> or the execution of a trigger. A transient table is identified by a <query name> if it arises during the evaluation of a <query expression>, or by a <transition table name> if it arises during the execution of a trigger. Such tables exist only for the duration of the executing SQL-statement containing the <query expression> or for the duration of the executing trigger.

A table is either *updatable* or *not updatable*. An updatable table has at least one *updatable column*. All base tables are updatable and every column of a base table is updatable. Derived tables and transient tables are either updatable or not updatable. The Syntax Rules of Subclause 7.6, “<table reference>”, determine whether a derived table *T* is updatable, and which of the columns of *T* are updatable. A view is *updatable* if the derived table that defines the view is updatable. A column of a view is *updatable* if the corresponding column of the derived table that defines the view is updatable. A view is *trigger updatable* if an update INSTEAD OF trigger is defined on that view. A view is *trigger deletable* if a delete INSTEAD OF trigger is defined on that view.

Some updatable tables, including all base tables whose row type is not derived from a user-defined type that is not instantiable, are also *insertable-into*. Transient tables are not insertable-into. The Syntax Rules of Subclause 7.6, “<table reference>”, determine whether a derived table *T* is insertable-into. A view is *insertable-into* if the derived table that defines the view is insertable-into. A view is *trigger insertable-into* if an insert INSTEAD OF trigger is defined on that view.

The operation of update on a table *T* is permitted if *T* is updatable or trigger updatable, subject to constraining Access Rules and Conformance Rules. The operation of delete on a table *T* is permitted if *T* is updatable or trigger deletable, subject to constraining Access Rules and Conformance Rules. The operation of insert on a table *T* is permitted if *T* is insertable-into or trigger insertable-into, subject to constraining Access Rules and Conformance Rules.

A *grouped table* is a set of groups derived during the evaluation of a <group by clause>. A group *G* is a collection of rows in which, for every grouping column *GC*, if the value of *GC* in some row is not distinct from *GV*, then the value of *GC* in every row is *GV*; moreover, if *R1* is a row in group *G1* of grouped table *GT* and *R2* is a row in *GT* such that for every grouping column *GC* the value of *GC* in *R1* is not distinct from the value of *GC* in *R2*, then *R2* is in *G1*. Every row in *GT* is in exactly one group. A group may be considered as a table. Set functions operate on groups.

A global temporary table is a named table defined by a <table definition> that specifies GLOBAL TEMPORARY. Global temporary tables are effectively materialized only when referenced in an SQL-session. Every SQL-session that references a global temporary table causes a distinct instance of that global temporary table (*i.e.*, a multiset of rows that is visible only to that SQL-session) to be materialized. That is, the multiset of rows that is reference by the <table name> of a global temporary table cannot be shared between SQL-sessions.

A created local temporary table is a named table defined by a <table definition> that specifies LOCAL TEMPORARY. Created local temporary tables are effectively materialized only when referenced in an SQL-session. Every SQL-client module in every SQL-session that references a created local temporary table causes a distinct instance of that created local temporary table (*i.e.*, a multiset of rows that is visible only to that SQL-client module during that SQL-session) to be materialized. That is, the multiset of rows that is referenced by the <table name> of a created local temporary table cannot be shared between SQL-sessions, nor between SQL-client modules that execute during an SQL-session.

The definition of a global temporary table or a created local temporary table appears in a schema. In SQL language, the name and the scope of the name of a global temporary table or a created local temporary table are indistinguishable from those of a persistent base table.

A declared local temporary table is a named table defined by a <temporary table declaration>.

An SQL-client module declared local temporary table is a declared local temporary table defined in an <SQL-client module definition>. An SQL-client module declared local temporary table is effectively materialized the first time any <externally-invoked procedure> in the <SQL-client module definition> that contains the <temporary table declaration> is executed, and it persists for that SQL-session. Every SQL-client module in every SQL-session that references an SQL-client module declared local temporary table causes a distinct instance of that declared local temporary table (*i.e.*, a multiset of rows that is visible only to that SQL-client module during that SQL-session) to be materialized. That is, the multiset of rows that is referenced by the <table name> of an SQL-client module declared local temporary table cannot be shared between SQL-sessions, nor between SQL-client modules that execute during an SQL-session.

All references to a declared local temporary table are prefixed by “MODULE.”.

The materialization of a temporary table does not persist beyond the end of the SQL-session in which the table was materialized. Temporary tables are effectively empty at the start of an SQL-session. Temporary tables that are created within an SQL-session *S1* are also effectively empty after creation. If an <as subquery clause> that

specifies WITH DATA is specified, the materialization of the temporary table within *SI* is initialized as specified by the General Rules of [Subclause 11.3](#), “<table definition>”.

The <table name> of a persistent base table, global temporary table, created local temporary table, or SQL-client module declared local temporary table is used to identify a particular multiset of rows, as follows:

- The <table name> of a persistent base table uniquely identifies a multiset of rows.
- The <table name> of a global temporary table, together with an SQL-session identifier, uniquely identifies a multiset of rows.
- The <table name> of a created local temporary table, together with an SQL-session identifier and an SQL-client module name, uniquely identifies a multiset of rows.
- The <table name> of an SQL-client module declared local temporary table, together with an SQL-session identifier and an SQL-client module name, uniquely identifies a multiset of rows.

4.14.3 Table descriptors

This Subclause is modified by [Subclause 4.10.2](#), “Table descriptors”, in ISO/IEC 9075-9.

[09] A table is described by a table descriptor. A table descriptor is either a transient table descriptor, a base table descriptor, a view descriptor, or a derived table descriptor (for a derived table that is not a view).

Every table descriptor includes:

- The column descriptor of each column in the table.
- The name, if any, of the structured type, if any, associated with the table.
- An indication of whether the table is insertable-into or not.
- An indication of whether the table is a referenceable table or not, and an indication of whether the self-referencing column is a system-generated, a user-generated, or a derived self-referencing column.
- A list, possibly empty, of the names of its direct supertables.
- A list, possibly empty, of the names of its direct subtables.

A transient table descriptor describes a transient table. In addition to the components of every table descriptor, a transient table descriptor includes:

- If the transient table is defined by a <with list element> contained in a <query expression>, then the <query name>. If the transient table is defined by a <trigger definition> then the <transition table name>.

A base table descriptor describes a base table. In addition to the components of every table descriptor, a base table descriptor includes:

- The name of the base table.
- An indication of whether the table is a persistent base table, a global temporary table, a created local temporary table, or a declared local temporary table.

- If the base table is a global temporary table, a created local temporary table, or a declared local temporary table, then an indication of whether ON COMMIT PRESERVE ROWS was specified or ON COMMIT DELETE ROWS was specified or implied.
- The descriptor of each table constraint specified for the table.
- A non-empty set of functional dependencies, according to the rules given in [Subclause 4.18, “Functional dependencies”](#).
- A non-empty set of candidate keys, according to the rules of [Subclause 4.19, “Candidate keys”](#).
- A preferred candidate key, which may or may not be additionally designated the primary key, according to the Rules in [Subclause 4.18, “Functional dependencies”](#).

A derived table descriptor describes a derived table. In addition to the components of every table descriptor, a derived table descriptor includes:

- The <query expression> that defines how the table is to be derived.
- An indication of whether the derived table is updatable or not.
- An indication of whether the derived table is simply updatable or not.

A view descriptor describes a view. In addition to the components of a derived table descriptor, a view descriptor includes:

- The name of the view.
- An indication of whether the view has the CHECK OPTION; if so, whether it is to be applied as CASCADED or LOCAL.
- The original <query expression> of the view.
- An indication of whether the view is trigger updatable.
- An indication of whether the view is trigger insertable-into.
- An indication of whether the view is trigger deletable.

4.14.4 Relationships between tables

This Subclause is modified by [Subclause 4.10.3, “Relationships between tables”](#), in ISO/IEC 9075-9.

The terms *simply underlying table*, *underlying table*, *leaf underlying table*, *generally underlying table*, and *leaf generally underlying table* define a relationship between a derived table or cursor and other tables.

The simply underlying tables of derived tables are defined in the Syntax Rules of [Subclause 7.12, “<query specification>”](#), and [Subclause 7.13, “<query expression>”](#). A <table or query name> has no simply underlying tables. The simply underlying table of a cursor is the <cursor specification> included in the cursor's result set descriptor (or, for standing cursors, the <cursor specification> included in the cursor declaration descriptor).

The underlying tables of a derived table or cursor are the simply underlying tables of the derived table or cursor and the underlying tables of the simply underlying tables of the derived table or cursor.

The leaf underlying tables of a derived table or cursor are the underlying tables of the derived table or cursor that do not themselves have any underlying tables.

09 The generally underlying tables of a derived table or cursor are the underlying tables of the derived table or cursor and, for each underlying table of the derived table or cursor that is a <table or query name> *TORQN*, the generally underlying tables of *TORQN*, which are defined as follows:

- 09 If *TORQN* identifies a base table or if *TORQN* is a <transition table name>, then *TORQN* has no generally underlying tables.
- If *TORQN* is a <query name>, then the generally underlying tables of *TORQN* are the <query expression body> *QEB* of the <with list element> identified by *TORQN* and the generally underlying tables of *QEB*.
- If *TORQN* identifies a view *V*, then the generally underlying tables of *TORQN* are the <query expression> *QEV* included in the view descriptor of *V* and the generally underlying tables of *QEV*.

The leaf generally underlying tables of a derived table or cursor are the generally underlying tables of the derived table or cursor that do not themselves have any generally underlying tables.

4.14.5 Referenceable tables, subtables, and supertables

A table *RT* whose row type is derived from a structured type *ST* is called a *typed table*. Only a base table or a view can be a typed table. A typed table has columns corresponding, in name and declared type, to every attribute of *ST* and one other column *REFC* that is the self-referencing column of *RT*; let *REFCN* be the <column name> of *REFC*. The declared type of *REFC* is necessarily *REF(ST)* and the nullability characteristic of *REFC* is *known not nullable*. If *RT* is a base table, then the table constraint “*UNIQUE(REFCN)*” is implicit in the definition of *RT*. A typed table is called a *referenceable table*. A self-referencing column cannot be updated. Its value is determined during the insertion of a row into the referenceable table. The value of a system-generated self-referencing column and a derived self-referencing column is automatically generated when the row is inserted into the referenceable table. The value of a user-generated self-referencing column is supplied as part of the candidate row to be inserted into the referenceable table.

A table T_a is a *direct subtable* of another table T_b if and only if the <table name> of T_b is contained in the <subtable clause> contained in the <table definition> or <view definition> of T_a . Both T_a and T_b shall be created on a structured type and the structured type of T_a shall be a direct subtype of the structured type of T_b .

A table T_a is a *subtable* of a table T_b if and only if any of the following are true:

- 1) T_a and T_b are the same named table.
- 2) T_a is a direct subtable of T_b .
- 3) There is a table T_c such that T_a is a direct subtable of T_c and T_c is a subtable of T_b .

A table *T* is considered to be one of its own subtables. Subtables of *T* other than *T* itself are called its *proper subtables*. A table shall not have itself as a proper subtable.

A table T_b is called a *supertable* of a table T_a if T_a is a subtable of T_b . If T_a is a direct subtable of T_b , then T_b is called a *direct supertable* of T_a . A table that is not a subtable of any other table is called a *maximal supertable*.

Let T_a be a maximal supertable and T be a subtable of T_a . The set of all subtables of T_a (which includes T_a itself) is called the *subtable family* of T or (equivalently) of T_a . Every subtable family has exactly one maximal supertable.

A *leaf table* is a table that does not have any proper subtables.

Those columns of a subtable T_a of a structured type ST_a that correspond to the inherited attributes of ST_a are called *inherited columns*. Those columns of T_a that correspond to the originally-defined attributes of ST_a are called *originally-defined columns*.

Let TB be a subtable of TA . Let SLA be the <value expression> sequence implied by the <select list> “*” in the <query specification> “SELECT * FROM TA ”. For every row RB in the value of TB there exists exactly one row RA in the value of TA such that RA is the result of the <row subquery> “SELECT SLA FROM VALUES RRB ”, where RRB is some <row value constructor> whose value is RB . RA is said to be the *superrow* in TA of RB and RB is said to be the *subrow* in TB of RA . If TA is a base table, then the one-to-one correspondence between superrows and subrows is guaranteed by the requirement for a unique constraint to be specified for some supertable of TA . If TA is a view, then such one-to-one correspondence is guaranteed by the requirement for a unique constraint to be specified on the leaf generally underlying table of TA .

Users shall have the UNDER privilege on a table before they can use the table in a subtable definition. A table can have more than one proper subtable. Similarly, a table can have more than one proper supertable.

4.14.6 Operations involving tables

Table values are operated on and returned by <query expression>s. The syntax of <query expression> includes various internal operators that operate on table values and return table values. In particular, every <query expression> effectively includes at least one <from clause>, which operates on one or more table values and returns a single table value. A table value operated on by a <from clause> is specified by a <table reference>.

An operation involving a table T may define a *range variable* RV that ranges over rows of T , referencing each row in turn in an implementation-dependent order. Thus, each reference to RV references exactly one row of T . T is said to be the *table associated with* RV .

In a <table reference>, ONLY can be specified to exclude from the result rows that have subrows in proper subtables of the referenced table.

In a <table reference>, <sample clause> can be specified to return a subset of result rows depending on the <sample method> and <sample percentage>. If the <sample clause> contains <repeatable clause>, then repeated executions of that <table reference> return a result table with identical rows for a given <repeat argument>, provided certain implementation-defined conditions are satisfied.

Certain table updating operations, specified by SQL-data change statements, are available in connection with updatable tables and trigger updatable tables, subject to applicable Access Rules and Conformance Rules. The value of an updatable table T or a trigger updatable table T is determined by the result of the mostly recently executed SQL-data change statement (see Subclause 4.33.2, “SQL-statements classified by function”) operating on T . An SQL-data change statement on a base table T or a derived table T has a *primary effect* (on T itself) and zero or more *secondary effects* (not necessarily on T). The effect of an SQL-data change statement on a viewed table is specified by the General Rules of Subclause 15.9, “Effect of deleting some rows from a viewed table”, Subclause 15.12, “Effect of inserting a table into a viewed table”, and Subclause 15.15, “Effect of replacing some rows in a viewed table”.

4.14 Tables

The primary effect of a <delete statement: positioned> on a base table *T* or a derived table *T* is to delete exactly one specified row from *T*. The primary effect of a <delete statement: searched> on a base table *T* or a derived table *T* is to delete zero or more rows from *T*.

The primary effect of a <truncate table statement> on a table *T* is to delete all rows from *T* without causing the execution of any trigger specified for *T*.

The primary effect of an <update statement: positioned> on a base table *T* or a derived table *T* is to replace exactly one specified row in *T* with some specified row. The primary effect of an <update statement: searched> on a base table *T* or a derived table *T* is to replace zero or more rows in *T*.

The primary effect of an <insert statement> on base table *T* or a derived table *T* is to insert into *T* each of the zero or more rows contained in a specified table.

The primary effect of a <merge statement> on base table *T* or a derived table *T* is to replace zero or more rows in *T* with specified rows and/or to insert into *T* zero or more specified rows.

Each of the table updating operations, when applied to a base table *T* or a derived table *T*, can have various secondary effects. Such secondary effects include alteration or reversal of the primary effect. Secondary effects might arise from the existence of:

- Underlying tables of *T*, other than *T* itself, whose values might be subject to secondary effects.
- Updatable views whose <view definition>s do not specify **WITH CASCADED CHECK OPTION**.
- Cascaded operations specified in connection with integrity constraints involving underlying tables of *T*, which might result in secondary effects on tables referenced by such constraints.
- Proper subtables and proper supertables of *T*, whose values might be affected by updating operations on *T*.
- Triggers specified for underlying tables of *T*, which might specify table updating operations on updatable tables other than *T*.

4.14.7 Identity columns

The columns of a base table *BT* can optionally include not more than one *identity column*. The declared type of an identity column is either an exact numeric type with scale 0 (zero), **INTEGER** for example, or a distinct type whose source type is an exact numeric type with scale 0 (zero). An identity column has a *start value*, an *increment*, a *maximum value*, a *minimum value*, and a *cycle option*. An identity column is associated with an internal sequence generator *SG*. Let *IC* be the identity column of *BT*. When a row *R* is presented for insertion into *BT*, if *R* does not contain a column corresponding to *IC*, then the value *V* for *IC* in the row inserted into *BT* is obtained by applying the General Rules of [Subclause 9.23](#), “[Generation of the next value of a sequence generator](#)”, to *SG*. The definition of an identity column may specify **GENERATED ALWAYS** or **GENERATED BY DEFAULT**.

NOTE 30 — “Start value”, “increment”, “maximum value”, “minimum value”, and “cycle option” are defined in [Subclause 4.21](#), “[Sequence generators](#)”.

NOTE 31 — The notion of an internal sequence generator being associated with an identity column is used only for definitional purposes in this International Standard.

4.14.8 Base columns and generated columns

A column of a base table is either a *base column* or a *generated column*. A base column is one that is not a generated column. A generated column is one whose values are determined by evaluation of a *generation expression*, a <value expression> whose declared type is by implication that of the column. A generation expression can reference base columns of the base table to which it belongs but cannot otherwise access SQL-data. Thus, the value of the field corresponding to a generated column in row *R* is determined by the values of zero or more other fields of *R*.

A generated column *GC* depends on each column that is referenced by a <column reference> in its generation expression, and each such referenced column is a *parametric column* of *GC*.

4.14.9 Windowed tables

A *windowed table* is a table together with one or more windows. A *window* is a transient data structure associated with a <table expression>. A window is defined explicitly by a <window definition> or implicitly by an <in-line window specification>. Implicitly defined windows have an implementation-dependent window name. A window is used to specify window partitions and window frames, which are collections of rows used in the definition of <window function>s.

Every window defines a *window partitioning* of the rows of the <table expression>. The window partitioning is specified by a list of columns. Window partitioning is similar to forming groups of a grouped table. However, unlike grouped tables, each row is retained in the result of the <table expression>. The *window partition* of a row *R* is the collection of rows *R2* that are not distinct from *R*, for all columns enumerated in the window partitioning clause. The window partitioning clause is optional; if omitted, there is a single window partition consisting of all the rows in the result.

If a <table expression> is grouped and also has a window, then there is a syntactic transformation that segregates the grouping into a <derived table>, so that the window partitions consist of rows of the <derived table> rather than groups of rows.

A window may define a *window ordering* of rows within each window partition defined by the window. The window ordering of rows within window partitions is specified by a list of <value expression>s, followed by ASC (for ascending order) or DESC (for descending order). In addition, NULLS FIRST or NULLS LAST may be specified, to indicate whether a null value should appear before or after all non-null values in the ordered sequence of each <value expression>.

Optionally, a window may define a *window frame* for each row *R*. A window frame is always defined relative to the current row. A window frame is specified by up to four syntactic elements:

- The choice of RANGE, to indicate a logical definition of the window frame by offsetting forward or backward from the current row by an increment or decrement to the sort key; or ROWS, to indicate a physical definition of the window frame, by counting rows forward or backward from the current row.
- A starting row, which may be the first row of the window partition of *R*, the current row, or some row determined by a logical or physical offset from the current row.
- An ending row, which may be the last row of the window partition of *R*, the current row, or some row determined by a logical or physical offset from the current row.
- A <window frame exclusion>, indicating whether to exclude the current row and/or its peers (if not already excluded by being prior to the starting row or after the ending row).

4.14 Tables

A window is described by a *window structure descriptor*, including:

- The window name.
- Optionally, the ordering window name—that is, the name of another window, called the *ordering window*, that is used to define the partitioning and ordering of the present window.
- The window partitioning clause—that is, a <window partition clause> if any is specified in either the present <window specification> or in the window descriptor of the ordering window.
- The window ordering clause—that is, a <window order clause> if any is specified in either the present <window specification> or in the window descriptor of the ordering window.
- The window framing clause—that is, a <window frame clause>, if any.

In general, two <window function>s are computed independently, each one performing its own sort of its data, even if they use the same data and the same <sort specification list>. Since sorts may specify partial orderings, the computation of <window function>s is inevitably non-deterministic to the extent that the ordering is not total. Nevertheless, the user may desire that two <window function>s be computed using the same ordering, so that, for example, two moving aggregates move through the rows of a partition in precisely the same order. Two <window function>s are computed using the same (possibly non-deterministic) window ordering of the rows if any of the following are true:

- The <window function>s identify the same window structure descriptor.
- The <window function>s' window structure descriptors have window partitioning clauses that enumerate the same number of column references, and those column references are pairwise equivalent in their order of occurrence; and their window structure descriptors have window ordering clauses with the same number of <sort key>s, and those <sort key>s are all column references, and those column references are pairwise equivalent in their order of occurrence, and the <sort specification>s pairwise specify or imply <collate clause>s that specify equivalent <collation name>s, the same <ordering specification> (ASC or DESC), and the same <null ordering> (NULLS FIRST or NULLS LAST).
- The window structure descriptor of one <window function> is the ordering window of the other <window function>, or both window structure descriptors identify the same ordering window.

4.15 Data analysis operations (involving tables)

This Subclause is modified by Subclause 4.4, “Data analysis operations (involving tables)”, in ISO/IEC 9075-14.

4.15.1 Introduction to data analysis operations

A data analysis function is a function that returns a value derived from a number of rows in the result of a <table expression>. A data analysis function may only be invoked as part of a <query specification> or <select statement: single row>, and then only in certain contexts, identified below. A data analysis function is one of:

- A group function, which is invoked on a grouped table and computes a grouping operation or an aggregate function from a group of the grouped table.
- A window function, which is invoked on a windowed table and computes a rank, row number or window aggregate function.

4.15.2 Group functions

A group function may only appear in the <select list>, <having clause> or <window clause> of a <query specification> or <select statement: single row>, or in the <order by clause> simply contained in a <query expression> that is a simple table query.

A group function is one of:

- The *grouping operation*.
- A *group aggregate function*.

The grouping operation is of the form `GROUPING(<column reference>)`. The result of such an invocation is 1 (one) in the case of a row whose values are the results of aggregation over that <column reference> during the execution of a grouped query containing CUBE, ROLLUP, or GROUPING SETS, and 0 (zero) otherwise.

4.15.3 Window functions

A window function is a function whose result for a given row is derived from the window frame of that row as defined by a window structure descriptor of a windowed table. Window functions may only appear in the <select list> of a <query specification> or <select statement: single row>, or the <order by clause> simply contained in a <query expression> that is a simple table query.

A window function is one of:

- A rank function.
- A distribution function.
- The row number function.
- A window aggregate function.

The rank functions compute the ordinal rank of a row *R* within the window partition of *R* as defined by a window structure descriptor, according to the window ordering of those rows, also specified by the same window structure descriptor. Rows that are not distinct with respect to the window ordering within their window partition are assigned the same rank. There are two variants, indicated by the keywords RANK and DENSE_RANK.

- If RANK is specified, then the rank of row *R* is defined as 1 (one) plus the number of rows that precede *R* and are not peers of *R*.

NOTE 32 — This implies that if two or more rows are not distinct with respect to the window ordering, then there will be one or more gaps in the sequential rank numbering.

- If DENSE_RANK is specified, then the rank of row *R* is defined as the number of rows preceding and including *R* that are distinct with respect to the window ordering.

NOTE 33 — This implies that there are no gaps in the sequential rank numbering of rows in each window partition.

The distribution functions compute a relative rank of a row *R* within the window partition of *R* defined by a window structure descriptor, expressed as an approximate numeric ratio between 0.0 and 1.0. There are two variants, indicated by the keywords PERCENT_RANK and CUME_DIST.

4.15 Data analysis operations (involving tables)

- If PERCENT_RANK is specified, then the relative rank of a row R is defined as $(RK-1)/(NR-1)$, where RK is defined to be the RANK of R and NR is defined to be the number of rows in the window partition of R .
- If CUME_DIST is specified, then the relative rank of a row R is defined as NP/NR , where NP is defined to be the number of rows preceding or peer with R in the window ordering of the window partition of R and NR is defined to be the number of rows in the window partition of R .

The ROW_NUMBER function computes the sequential row number, starting with 1 (one) for the first row, of the row within its window partition according to the window ordering of the window.

The window aggregate functions compute an aggregate value (COUNT, SUM, AVG, *etc.*), the same as a group aggregate function, except that the computation aggregates over the window frame of a row rather than over a group of a grouped table. The hypothetical set functions are not permitted as window aggregate functions.

4.15.4 Aggregate functions

This Subclause is modified by Subclause 4.4.1, “Aggregate functions”, in ISO/IEC 9075-14.

An aggregate function is a function whose result is derived from an aggregation of rows defined by one of:

- The grouping of a grouped table, in which case the aggregate function is a group aggregate function, or set function, and for each group there is one aggregation, which includes every row in the group.
- The window frame of a row R of a windowed table relative to a particular window structure descriptor, in which case the aggregate function is a window aggregate function, and the aggregation consists of every row in the window frame of R , as defined by the window structure descriptor.

Optionally, the collection of rows in an aggregation may be filtered, retaining only those rows that satisfy a <search condition> that is specified by a <filter clause>.

The result of the aggregate function COUNT (*) is the number of rows in the aggregation.

Every other aggregate function may be classified as a *unary group aggregate function*, a *binary group aggregate functions*, an *inverse distribution*, or a *hypothetical set function*.

Every unary aggregate function takes an arbitrary <value expression> as the argument; most unary aggregate functions can optionally be qualified with either DISTINCT or ALL. Of the rows in the aggregation, the following do not qualify:

- If DISTINCT is specified, then redundant duplicates.
- Every row in which the <value expression> evaluates to the null value.

If no row qualifies, then the result of COUNT is 0 (zero), and the result of any other aggregate function is the null value.

14 Otherwise (*i.e.*, at least one row qualifies), the result of the aggregate function is:

- If COUNT <value expression> is specified, then the number of rows that qualify.
- If SUM is specified, then the sum of <value expression> evaluated for each row that qualifies.
- If AVG is specified, then the average of <value expression> evaluated for each row that qualifies.

4.15 Data analysis operations (involving tables)

- If MAX is specified, then the maximum value of <value expression> evaluated for each row that qualifies.
- If MIN is specified, then the minimum value of <value expression> evaluated for each row that qualifies.
- If EVERY is specified, then True if the <value expression> evaluates to True for every row that qualifies; otherwise, False.
- If ANY or SOME is specified, then True if the <value expression> evaluates to True for at least one row remaining in the group; otherwise, False.
- If VAR_POP is specified, then the population variance of <value expression> evaluated for each row remaining in the group, defined as the sum of squares of the difference of <value expression> from the mean of <value expression>, divided by the number of rows remaining.
- If VAR_SAMP is specified, then the sample variance of <value expression> evaluated for each row remaining in the group, defined as the sum of squares of the difference of <value expression> from the mean of <value expression>, divided by the number of rows remaining minus 1 (one).
- If STDDEV_POP is specified, then the population standard deviation of <value expression> evaluated for each row remaining in the group, defined as the square root of the population variance.
- If STDDEV_SAMP is specified, then the sample standard deviation of <value expression> evaluated for each row remaining in the group, defined as the square root of the sample variance.
- If ARRAY_AGG is specified, then an array value with one element formed from the <value expression> evaluated for each row that qualifies.

Neither DISTINCT nor ALL are allowed to be specified for VAR_POP, VAR_SAMP, STDDEV_POP, or STDDEV_SAMP; redundant duplicates are not removed when computing these functions.

The binary aggregate functions take a pair of arguments, the <dependent variable expression> and the <independent variable expression>, which are both <numeric value expression>s. Any row in which either argument evaluates to the null value is removed from the group. If there are no rows remaining in the group, then the result of REGR_COUNT is 0 (zero), and the other binary aggregate functions result in the null value. Otherwise, the computation concludes and the result is:

- If REGR_COUNT is specified, then the number of rows remaining in the group.
- If COVAR_POP is specified, then the population covariance, defined as the sum of products of the difference of <independent variable expression> from its mean times the difference of <dependent variable expression> from its mean, divided by the number of rows remaining.
- If COVAR_SAMP is specified, then the sample covariance, defined as the sum of products of the difference of <independent variable expression> from its mean times the difference of <dependent variable expression> from its mean, divided by the number of rows remaining minus 1 (one).
- If CORR is specified, then the correlation coefficient, defined as the ratio of the population covariance divided by the product of the population standard deviation of <independent variable expression> and the population standard deviation of <dependent variable expression>.
- If REGR_R2 is specified, then the square of the correlation coefficient.
- If REGR_SLOPE is specified, then the slope of the least-squares-fit linear equation determined by the (<independent variable expression>, <dependent variable expression>) pairs.
- If REGR_INTERCEPT is specified, then the y-intercept of the least-squares-fit linear equation determined by the (<independent variable expression>, <dependent variable expression>) pairs.

4.15 Data analysis operations (involving tables)

- If REGR_SXX is specified, then the sum of squares of <independent variable expression>.
- If REGR_SYY is specified, then the sum of squares of <dependent variable expression>.
- If REGR_SXY is specified, then the sum of products of <independent variable expression> times <dependent variable expression>.
- If REGR_AVGX is specified, then the average of <independent variable expression>.
- If REGR_AVGY is specified, then the average of <dependent variable expression>.

There are two inverse distribution functions, PERCENTILE_CONT and PERCENTILE_DISC. Both inverse distribution functions specify an argument and an ordering of a value expression. The value of the argument should be between 0 (zero) and 1 (one) inclusive. The value expression is evaluated for each row of the group, nulls are discarded, and the remaining rows are ordered. The computation concludes:

- If PERCENTILE_CONT is specified, by considering the pair of consecutive rows that are indicated by the argument, treated as a fraction of the total number of rows in the group, and interpolating the value of the value expression evaluated for these rows.
- If PERCENTILE_DISC is specified, by treating the group as a window partition of the CUME_DIST window function, using the specified ordering of the value expression as the window ordering, and returning the first value expression whose cumulative distribution value is greater than or equal to the argument.

The hypothetical set functions are related to the window functions RANK, DENSE_RANK, PERCENT_RANK, and CUME_DIST, and use the same names, though with a different syntax. These functions take an argument *A* and an ordering of a value expression *VE*. *VE* is evaluated for all rows of the group. This collection of values is augmented with *A*; the resulting collection is treated as a window partition of the corresponding window function whose window ordering is the ordering of the value expression. The result of the hypothetical set function is the value of the eponymous window function for the hypothetical “row” that contributes *A* to the collection.

4.16 Determinism

In general, an operation is *deterministic* if that operation assuredly computes identical results when repeated with identical input values. For an SQL-invoked routine, the values in the argument list are regarded as the input; otherwise, the SQL-data and the set of privileges by which they are accessed is regarded as the input. Differences in the ordering of rows, as permitted by General Rules that specify implementation-dependent behavior, are not regarded as significant to the question of determinism.

NOTE 34 — Transaction isolation levels have a significant impact on determinism, particularly isolation levels other than SERIALIZABLE. However, this International Standard does not address that impact, particularly because of the difficulty in clearly specifying that impact without appearing to mandate implementation techniques (such as row or page locking) and because different SQL-implementations almost certainly resolve the issue in significantly different ways.

Recognizing that an operation is deterministic is a difficult task, it is in general not mandated by this International Standard. SQL-invoked routines are regarded as deterministic if the routine is declared to be DETERMINISTIC; that is, the SQL-implementation trusts the definer of the SQL-invoked routine to correctly declare that the routine is deterministic. For other operations, this International Standard does not label an operation as deterministic; instead it identifies certain operations as “possibly non-deterministic”. Specific definitions can be found in other subclauses relative to <value expression>, <table reference>, <table primary>, <query specification>, <query expression>, and <SQL procedure statement>.

Certain <boolean value expression>s are identified as “retrospectively deterministic”. A retrospectively deterministic <boolean value expression> has the property that if it is *True* at one point time, then it is *True* for all later points in time if re-evaluated for the identical SQL-data by an arbitrary user with the identical set of privileges. The precise definition is found in Subclause 6.34, “<boolean value expression>”.

4.17 Integrity constraints

4.17.1 Overview of integrity constraints

Integrity constraints, generally referred to simply as constraints, define the valid states of SQL-data by constraining the values in the base tables. A constraint is described by a constraint descriptor. A constraint is either a table constraint, a domain constraint, or an assertion and is described by, respectively, a table constraint descriptor, a domain constraint descriptor, or an assertion descriptor. Every constraint descriptor includes:

- The name of the constraint.
- An applicable <search condition>.

NOTE 35 — The applicable <search condition> included in the descriptor is not necessarily the <search condition> that might be contained in the SQL-statement whose execution brings the constraint descriptor into existence. The General Rules for the SQL-statement in question specify the applicable <search condition> to be included in the constraint descriptor, in some cases deriving it from a given <search condition>. For example, the syntax for table constraints allows universal quantification over the rows of the table in question to be implicit; in the applicable <search condition> included in the descriptor, that universal quantification is made explicit, to allow for uniform treatment of all types of constraint.

- An indication of whether or not the constraint is deferrable.
- An indication of whether the initial constraint mode is *deferred* or *immediate*.

No integrity constraint shall be defined using a <search condition> that is not retrospectively deterministic.

4.17.2 Checking of constraints

Every constraint is either *deferrable* or *non-deferrable*. Within an SQL-transaction, every constraint has a constraint mode; if a constraint is non-deferrable, then its constraint mode is always *immediate*; otherwise, it is either *immediate* or *deferred*. Every constraint has an initial constraint mode that specifies the constraint mode for that constraint at the start of each SQL-transaction and immediately after definition of that constraint. If a constraint is deferrable, then its constraint mode may be changed (from *immediate* to *deferred*, or from *deferred* to *immediate*) by execution of a <set constraints mode statement>.

The checking of a constraint depends on its constraint mode within the current SQL-transaction. Whenever an SQL-statement is executed, every constraint whose mode is *immediate* is checked, at a certain point after any changes to SQL-data and schemas resulting from that execution have been effected, to see if it is satisfied. A constraint is *satisfied* if and only if the applicable <search condition> included in its descriptor evaluates to *True* or *Unknown*. If any constraint is not satisfied, then any changes to SQL-data or schemas resulting from executing that statement are canceled. (See the General Rules of Subclause 13.5, “<SQL procedure statement>”.

NOTE 36 — This includes SQL-statements that are executed as a direct result or an indirect result of executing a different SQL-statement. It also includes statements whose effects explicitly or implicitly include setting the constraint mode to *immediate*.

The constraint mode can be set to immediate either explicitly by execution of a <set constraints mode statement>, or implicitly at the end of the current SQL-transaction.

When a <commit statement> is executed, all constraints are effectively checked and, if any constraint is not satisfied, then an exception condition is raised and the SQL-transaction is terminated by an implicit <rollback statement>.

In contrast to constraint checking, the referential actions specified by a referential constraint are never deferred.

4.17.3 Table constraints

A *table constraint* is a constraint whose descriptor is included in a base table descriptor. Being associated with a particular base table allows for convenient syntactic shorthands in which universal quantification over the rows of the table in question is implied.

A table constraint is either a unique constraint, a referential constraint or a table check constraint. A table constraint is described by a table constraint descriptor, which is either a unique constraint descriptor, a referential constraint descriptor, or a table check constraint descriptor, respectively.

Every table constraint specified for base table *T* is implicitly a constraint on every subtable of *T*, by virtue of the fact that every row in a subtable is considered to have a corresponding superrow in every one of its supertables.

A unique constraint is satisfied if and only if no two rows in a table have the same non-null values in the *unique columns*. In addition, if the unique constraint was defined with PRIMARY KEY, then it requires that none of the values in the specified column or columns be a null value.

In addition to the components of every table constraint descriptor, a unique constraint descriptor includes:

- An indication of whether it was defined with PRIMARY KEY or UNIQUE.
- The names and positions of the *unique columns* specified in the <unique column list>.

If the table descriptor for base table *T* includes a unique constraint descriptor indicating that the unique constraint was defined with PRIMARY KEY, then the columns of that unique constraint constitute the *primary key* of *T*. A table that has a primary key cannot have a proper supertable.

In addition to the components of every table constraint descriptor, a referential constraint descriptor includes:

- A list of the names of the *referencing columns* specified in the <referencing columns>.
- The *referenced table* specified in the <referenced table and columns>.
- A list of the names of the *referenced columns* specified in the <referenced table and columns>.
- The <match type> and the <referential triggered action>.

The choices for <match type> are MATCH SIMPLE, MATCH PARTIAL, and MATCH FULL; MATCH SIMPLE is the default. There is no semantic difference between these choices if there is only one referencing column (and hence, only one referenced column). There is also no semantic difference if all referencing columns are not nullable. If there is more than one referencing column, at least one of which is nullable, then the various <match type>s have the following semantics:

- MATCH SIMPLE: if at least one referencing column is null, then the row of the referencing table passes the constraint check. If all referencing columns are not null, then the row passes the constraint check if and only if there is a row of the referenced table that matches all the referencing columns.
- MATCH PARTIAL: if all referencing columns are null, then the row of the referencing table passes the constraint check. If at least one referencing column is not null, then the row passes the constraint check if and only if there is a row of the referenced table that matches all the non-null referencing columns.
- MATCH FULL: if all referencing columns are null, then the row of the referencing table passes the constraint check. If all referencing columns are not null, then the row passes the constraint check if and only if there is a row of the referenced table that matches all the referencing columns. If some referencing column is null and another referencing column is nonnull, then the row of the referencing table violates the constraint check.

The ordering of the lists of referencing column names and referenced column names is implementation-defined, but shall be such that corresponding column names occupy corresponding positions in each list.

In the case that a table constraint is a referential constraint, the table is referred to as the *referencing table*. The *referenced columns* of a referential constraint shall be the *unique columns* of some unique constraint of the *referenced table*.

A referential constraint is satisfied if one of the following conditions is true, depending on the <match type> specified in the <referential constraint definition>:

- If no <match type> was specified then, for each row *R1* of the *referencing table*, either at least one of the values of the *referencing columns* in *R1* shall be a null value, or the value of each referencing column in *R1* shall be equal to the value of the corresponding *referenced column* in some row of the *referenced table*.
- If MATCH FULL was specified then, for each row *R1* of the *referencing table*, either the value of every *referencing column* in *R1* shall be a null value, or the value of every *referencing column* in *R1* shall not be null and there shall be some row *R2* of the *referenced table* such that the value of each *referencing column* in *R1* is equal to the value of the corresponding *referenced column* in *R2*.
- If MATCH PARTIAL was specified then, for each row *R1* of the *referencing table*, there shall be some row *R2* of the *referenced table* such that the value of each *referencing column* in *R1* is either null or is equal to the value of the corresponding *referenced column* in *R2*.

The referencing table may be the same table as the referenced table.

A table check constraint is satisfied if and only if the specified <search condition> evaluates to *True* or *Unknown* for every row of the table to which it applies.

NOTE 37 — Consequently, an empty table satisfies every table check constraint that applies to it.

An <update rule> that does not contain NO ACTION specifies a *referential update action*. A <delete rule> that does not specify NO ACTION specifies a *referential delete action*. Referential update actions and referential delete actions are collectively called *referential actions*. Referential actions are carried out before, and are not part of, the checking of a referential constraint. Deferring a referential constraint defers the checking of the <search condition> of the constraint (a <match predicate>) but does not defer the referential actions of the referential constraint.

NOTE 38 — For example, if a referential update action such as ON UPDATE CASCADE is specified, then any UPDATE operation on the referenced table will be cascaded to the referencing table as part of the UPDATE operation, even if the referential constraint is deferred. Consequently, the referential constraint cannot become violated by the UPDATE statement. On the other hand, ON UPDATE SET DEFAULT could result in a violation of the referential constraint if there is no matching row after the referencing column is set to its default value. In addition, INSERT and UPDATE operations on the referencing table do not entail any automatic

enforcement of the referential constraint. Any such violations of the constraint will be detected when the referential constraint is eventually checked, at or before a commit.

The options for <update rule> are:

- **ON UPDATE CASCADE:** any change to a referenced column in the referenced table causes the same change to the corresponding referencing column in matching rows of the referencing table.
- **ON UPDATE SET NULL:** any change to a referenced column in the referenced table causes the corresponding referencing column in matching rows of the referencing table to be set to null.
- **ON UPDATE SET DEFAULT:** any change to a referenced column in the referenced table causes the corresponding referencing column in matching rows of the referencing table to be set to its default value.
- **ON UPDATE RESTRICT:** any change to a referenced column in the referenced table is prohibited if there is a matching row.
- **ON UPDATE NO ACTION** (the default): there is no referential update action; the referential constraint only specifies a constraint check.

NOTE 39 — Even if constraint checking is not deferred, **ON UPDATE RESTRICT** is a stricter condition than **ON UPDATE NO ACTION**. **ON UPDATE RESTRICT** prohibits an update to a particular row if there are any matching rows; **ON UPDATE NO ACTION** does not perform its constraint check until the entire set of rows to be updated has been processed.

The options for <delete rule> are:

- **ON DELETE CASCADE:** if a row of the referenced table is deleted, then all matching rows in the referencing table are deleted.
- **ON DELETE SET NULL:** if a row of the referenced table is deleted, then all referencing columns in all matching rows of the referencing table to be set to null.
- **ON DELETE SET DEFAULT:** if a row of the referenced table is deleted, then all referencing columns in all matching rows of the referencing table to be set to the column's default value.
- **ON DELETE RESTRICT:** it is prohibited to delete a row of the referenced table if that row has any matching rows in the referencing table.
- **ON DELETE NO ACTION** (the default): there is no referential delete action; the referential constraint only specifies a constraint check.

NOTE 40 — Even if constraint checking is not deferred, **ON DELETE RESTRICT** is a stricter condition than **ON DELETE NO ACTION**. **ON DELETE RESTRICT** prohibits a delete of a particular row if there are any matching rows; **ON DELETE NO ACTION** does not perform its constraint check until the entire set of rows to be deleted has been processed.

4.17.4 Domain constraints

A domain constraint is a constraint that is specified for a domain. It is applied to all columns that are based on that domain, and to all values cast to that domain.

A domain constraint is described by a domain constraint descriptor. In addition to the components of every constraint descriptor, a domain constraint descriptor includes:

- The template <search condition> for the generation of domain constraint usage <search condition>s.
- A possibly empty set of domain constraint usages.

A domain constraint usage descriptor is created implicitly by the evaluation of a <column definition> whose <data type or domain name> is a <domain name>. If C is such a column and D is the domain identified by the <domain name>, then every domain constraint DC defined for D implies a domain constraint usage, to the effect that each value in C satisfies DC .

In addition to the components of every table constraint descriptor, a domain constraint usage descriptor includes:

- The name of the applicable column.
- The applicable <search condition> that evaluates whether each value in C satisfies DC .

A domain constraint is satisfied by SQL-data if and only if, for every table T that has a column named C based on that domain, the applicable <search condition> recorded in the appropriate domain constraint usage evaluates to *True* or *Unknown*.

A domain constraint is satisfied by the result of a <cast specification> if and only if the specified template <search condition>, with each occurrence of the <general value specification> VALUE replaced by that result, evaluates to *True* or *Unknown*.

4.17.5 Assertions

An assertion is a constraint whose descriptor is an independent schema component not included in any table descriptor.

4.18 Functional dependencies

This Subclause is modified by Subclause 4.11, “Functional dependencies”, in ISO/IEC 9075-9.

4.18.1 Overview of functional dependency rules and notations

This Subclause is modified by Subclause 4.11.1, “Overview of functional dependency rules and notations”, in ISO/IEC 9075-9.

09 This Subclause defines *functional dependency* and specifies a minimal set of rules that a conforming implementation shall follow to determine functional dependencies and candidate keys in base tables and <query expression>s.

The rules in this Subclause may be freely augmented by implementation-defined rules, where indicated in this Subclause.

Let T be any table. Let CT be the set comprising all the columns of T , and let A and B be arbitrary subsets of CT , not necessarily disjoint and possibly empty.

Let “ $T: A \mapsto B$ ” (read “in T , A determines B ” or “ B is functionally dependent on A in T ”) denote the functional dependency of B on A in T , which is true if, for any possible value of T , any two rows that are not distinct for every column in A also are not distinct for every column in B . When the table T is understood from context, the abbreviation “ $A \mapsto B$ ” may also be used.

4.18 Functional dependencies

If $X \mapsto Y$ is some functional dependency in some table T , then X is a *determinant* of Y in T .

Let $A \mapsto B$ and $C \mapsto D$ be any two functional dependencies in T . The following are also functional dependencies in T :

— $A \text{ UNION } (C \text{ DIFFERENCE } B) \mapsto B \text{ UNION } D$

— $C \text{ UNION } (A \text{ DIFFERENCE } D) \mapsto B \text{ UNION } D$

NOTE 41 — Here, “UNION” denotes set union and “DIFFERENCE” denotes set difference.

These two rules are called the *rules of deduction* for functional dependencies.

Every table has an associated non-empty set of functional dependencies.

The set of functional dependencies is non-empty because $X \mapsto X$ for any X . A functional dependency of this form is an axiomatic functional dependency, as is $X \mapsto Y$ where Y is a subset of X . $X \mapsto Y$ is a non-axiomatic functional dependency if Y is not a subset of X .

4.18.2 General rules and definitions

In the following Subclauses, let a column $C1$ be a *counterpart* of a column $C2$ under qualifying table QT if $C1$ is specified by a column reference (or by a <value expression> that is a column reference) that references $C2$ and QT is the qualifying table of $C2$. If $C1$ is a counterpart of $C2$ under qualifying table $QT1$ and $C2$ is a counterpart of $C3$ under qualifying table $QT2$, then $C1$ is a counterpart of $C3$ under $QT2$.

The notion of counterparts naturally generalizes to sets of columns, as follows: If $S1$ and $S2$ are sets of columns, and there is a one-to-one correspondence between $S1$ and $S2$ such that each element of $S1$ is a counterpart of the corresponding element of $S2$, then $S1$ is a counterpart of $S2$.

The next Subclauses recursively define the notion of *known functional dependency*. This is a ternary relationship between a table and two sets of columns of that table. This relationship expresses that a functional dependency in the table is known to the SQL-implementation. All axiomatic functional dependencies are known functional dependencies. In addition, any functional dependency that can be deduced from known functional dependencies using the rules of deduction for functional dependency is a known functional dependency.

The next Subclauses also recursively define the notion of a “*BUC-set*”, which is a set of columns of a table (as in “ S is BUC-set”, where S is a set of columns).

NOTE 42 — “BUC” is an acronym for “base table unique constraint”, since the starting point of the recursion is a set of known not null columns comprising a nondeferrable unique constraint of a base table.

The notion of BUC-set is closed under the following deduction rule for BUC-sets: If $S1$ and $S2$ are sets of columns, $S1$ is a subset of $S2$, $S1 \mapsto S2$, and $S2$ is a BUC-set, then $S1$ is also a BUC-set.

NOTE 43 — A BUC-set may be empty, in which case there is at most one row in the table. This case shall be distinguished from a table with no BUC-set.

An SQL-implementation may define additional rules for determining BUC-sets, provided that every BUC-set S of columns of a table T shall have an associated base table BT such that every column of S has a counterpart in BT , and for any possible value of the columns of S , there is at most one row in BT having those values in those columns.

The next Subclauses also recursively define the notion of a “*BPK-set*”, which is a set of columns of a table (as in “ S is a BPK-set”, where S is a set of columns). Every BPK-set is a BUC-set.

NOTE 44 — “BPK” is an acronym for “base table primary key”, since the starting point of the recursion is a set of known not null columns comprising a nondeferrable primary key constraint of a base table.

The notion of BPK-set is closed under the following deduction rule for BPK-sets: If $S1$ and $S2$ are sets of columns, $S1$ is a subset of $S2$, $S1 \mapsto S2$, and $S2$ is a BPK-set, then $S1$ is also a BPK-set.

NOTE 45 — Like BUC-sets, a BPK-set may be empty.

An SQL-implementation may define additional rules for determining BPK-sets, provided that every BPK-set S is a BUC-set, and every member of S has a counterpart to a column in a primary key in the associated base table BT .

All applicable syntactic transformations (for example, to remove *, CUBE, or ROLLUP) shall be applied before using the rules to determine known functional dependencies, BUC-sets, and BPK-sets.

The following Subclauses use the notion of *AND-component* of a <search condition> SC , which is defined recursively as follows:

- If SC is a <boolean test> BT , then the only AND-component of SC is BT .
- If SC is a <boolean factor> BF , then the only AND-component of SC is BF .
- If SC is a <boolean term> of the form “ P AND Q ”, then the AND-components of SC are the AND-components of P and the AND-components of Q .
- If SC is a <boolean value expression> BVE that specifies OR, then the only AND-component of SC is BVE .

Let AC be an AND-component of SC such that AC is a <comparison predicate> whose <comp op> is <equals operator>. Let $RVE1$ and $RVE2$ be the two <row value predicand>s that are the operands of AC . Suppose that both $RVE1$ and $RVE2$ are <row value constructor predicand>s. Let n be the degree of $RVE1$. Let $RVEC1_i$ and $RVEC2_i$, $1 \text{ (one)} \leq i \leq n$, be the i -th <common value expression>, <boolean predicand>, or <row value constructor element> of $RVE1$ and $RVE2$, respectively. The <comparison predicate> “ $RVEC1_i = RVEC2_i$ ” is called an *equality AND-component* of SC .

4.18.3 Known functional dependencies in a base table

Let T be a base table and let CT be the set comprising all the columns of T .

A set of columns $S1$ of T is a *BPK-set* if it is the set of columns enumerated in some unique constraint UC of T , UC specifies PRIMARY KEY, and UC is nondeferrable.

A set of columns $S1$ of T is a *BUC-set* if it is the set of columns enumerated in some unique constraint UC of T , UC is nondeferrable, and every member of $S1$ is known not null.

If UCL is a set of columns of T such that UCL is a BUC-set, then $UCL \mapsto CT$ is a *known functional dependency* in T .

If GC is a generated column of T , then $D \mapsto GC$, where D is the set of parametric columns of GC , is a *known functional dependency* in T .

Implementation-defined rules may determine other known functional dependencies in T .

4.18.4 Known functional dependencies in a transition table

Let TT be the transition table defined in a trigger TR and let T be the subject table of TR . If TT is an old transition table or if TR is an AFTER trigger and TT is a new transition table, then the BPK-sets, BUC-sets, and known functional dependencies of TT are the same as the BPK-sets, BUC-sets, and known functional dependencies of T . If TR is a BEFORE trigger and TT is a new transition table, then no set of columns of TT is a BPK-set or a BUC-set and the known functional dependencies of TT are the axiomatic functional dependencies.

4.18.5 Known functional dependencies in <table value constructor>

Let R be the result of a <table value constructor>, and let CR be the set comprising all the columns of R .

No set of columns of R is a BPK-set or a BUC-set, except as determined by implementation-defined rules.

All axiomatic functional dependencies are *known functional dependencies* of a <table value constructor>. In addition, there may be implementation-defined known functional dependencies (for example, by examining the actual value of the <table value constructor>).

4.18.6 Known functional dependencies in a <joined table>

Let $T1$ and $T2$ denote the tables identified by the first and second <table reference>s of some <joined table> JT . Let R denote the table that is the result of JT . Let CT be the set of columns of the result of JT .

Every column of R has some counterpart in either $T1$ or $T2$. If NATURAL is specified or the <join specification> is a <named columns join>, then some columns of R may have counterparts in both $T1$ and $T2$.

A set of columns S of R is a *BPK-set* if S has some counterpart in $T1$ or $T2$ that is a BPK-set, every member of S is known not null, and $S \mapsto CT$ is a *known functional dependency* of R .

A set of columns S of R is a *BUC-set* if S has some counterpart in $T1$ or $T2$ that is a BUC-set, every member of S is known not null, and $S \mapsto CT$ is a *known functional dependency* of R .

NOTE 46 — The following rules for known functional dependencies in a <joined table> are not mutually exclusive. The set of known functional dependencies is the union of those dependencies generated by all applicable rules, including the rules of deduction presented earlier.

If $A \mapsto B$ is a known functional dependency in $T1$, CA is the counterpart of A in R , and CB is the counterpart of B in R , then $CA \mapsto CB$ is a known functional dependency in R when one of the following is true:

- CROSS, INNER, or LEFT is specified.
- RIGHT or FULL is specified and at least one of the following is true:
 - At least one column in A is known not nullable.
 - All columns in A are join partitioning columns.

If $A \mapsto B$ is a known functional dependency in $T2$, CA is the counterpart of A in R , and CB is the counterpart of B in R , then $CA \mapsto CB$ is a known functional dependency in R when one of the following is true:

- CROSS, INNER, or RIGHT is specified.

— LEFT or FULL is specified and at least one of the following is true:

- At least one column in A is known not nullable.
- All columns in A are join partitioning columns.

If $\langle \text{join condition} \rangle$ is specified, AP is an equality AND-component of the $\langle \text{search condition} \rangle$, one comparand of AP is a column reference CR , and the other comparand of AP is a $\langle \text{literal} \rangle$, then let CRC be the counterparts of CR in R . Let $\{\}$ denote the empty set. $\{\} \mapsto \{CRC\}$ is a *known functional dependency* in R if any of the following conditions is true:

- INNER is specified.
- If LEFT is specified and CR is a column reference to a column in $T1$.
- If RIGHT is specified and CR is a column reference to a column in $T2$.

NOTE 47 — An SQL-implementation may also choose to recognize $\{\} \mapsto \{CRC\}$ as a known functional dependency if the other comparand is a deterministic expression containing no column references.

If $\langle \text{join condition} \rangle$ is specified, AP is an equality AND-component of the $\langle \text{search condition} \rangle$, one comparand of AP is a column reference CRA , and the other comparand of AP is a column references CRB , then let $CRAC$ and $CRBC$ be the counterparts of CRA and CRB in R . $\{CRAC\} \mapsto \{CRBC\}$ is a *known functional dependency* in R if any of the following conditions is true:

- INNER is specified.
- If LEFT is specified and CRA is a column reference to a column in $T1$.
- If RIGHT is specified and CRA is a column reference to a column in $T2$.

NOTE 48 — An SQL-implementation may also choose to recognize the following as known functional dependencies: $\{CRAC\} \mapsto \{CRBC\}$ if CRA is known not nullable, CRA is a column of $T1$, and RIGHT or FULL is specified; or if CRA is known not nullable, CRA is a column of $T2$, and LEFT or FULL is specified.

NOTE 49 — An SQL-implementation may also choose to recognize similar known functional dependencies of the form $\{CRA_1, \dots, CRA_N\} \mapsto \{CRBC\}$ in case one comparand is a deterministic expression of column references CRA_1, \dots, CRA_N under similar conditions.

If NATURAL is specified, or if a $\langle \text{join specification} \rangle$ immediately containing a $\langle \text{named columns join} \rangle$ is specified, then let C_1, \dots, C_N be the column names of the corresponding join columns, for i between 1 (one) and N . Let SC be the $\langle \text{search condition} \rangle$:

```
( TN1.C1 = TN2.C1 )
AND
...
AND
( TN1.CN = TN2.CN )
```

Let $SLCC$ and SL be the $\langle \text{select list} \rangle$ s defined in the Syntax Rules of Subclause 7.7, “ $\langle \text{joined table} \rangle$ ”. Let JT be the $\langle \text{join type} \rangle$. Let $TN1$ and $TN2$ be the exposed $\langle \text{table or query name} \rangle$ or $\langle \text{correlation name} \rangle$ of tables $T1$ and $T2$, respectively. Let IR be the result of the $\langle \text{query expression} \rangle$:

```
SELECT SLCC, TN1.*, TN2.*
FROM TN1 JT JOIN TN2
ON SC
```

The following are recognized as additional *known functional dependencies* of IR :

4.18 Functional dependencies

- If INNER or LEFT is specified, then $\{ \text{COALESCE} (TN1.C_i, TN2.C_i) \} \mapsto \{ TN1.C_i \}$, for all i between 1 (one) and N .
- If INNER or RIGHT is specified, then $\{ \text{COALESCE} (TN1.C_i, TN2.C_i) \} \mapsto \{ TN2.C_i \}$, for all i between 1 (one) and N .

The *known functional dependencies* of R are the known functional dependencies of:

```
SELECT SL FROM IR
```

4.18.7 Known functional dependencies in a <table primary>

Let R be the result of some <table primary> TP . The BPK-sets, BUC-sets, and functional dependencies of R are determined as follows.

Case:

- If TP immediately contains a <table or query name> TQN (with or without ONLY), then the counterparts of the BPK-sets and BUC-sets of TQN are the BPK-sets and BUC-sets, respectively, of R . If $A \mapsto B$ is a functional dependency in the result of TQN , and AC and BC are the counterparts of A and B , respectively, then $AC \mapsto BC$ is a *known functional dependency* in R .
- If TP immediately contains a <derived table> DT , then the counterparts of the BPK-sets and BUC-sets of DT are the BPK-sets and BUC-sets, respectively, of R . If $A \mapsto B$ is a functional dependency in the result of DT , and AC and BC are the counterparts of A and B , respectively, then $AC \mapsto BC$ is a *known functional dependency* in R .
- If TP immediately contains a <lateral derived table> LDT , then the counterparts of the BPK-sets and BUC-sets of LDT are the BPK-sets and BUC-sets, respectively, of R . If $A \mapsto B$ is a functional dependency in the result of LDT , and AC and BC are the counterparts of A and B , respectively, then $AC \mapsto BC$ is a *known functional dependency* in R .
- If TP immediately contains a <collection derived table> CDT , and WITH ORDINALITY is specified, then let OC be the column name of the ordinality column and let CT be the set comprising all of the column names of the columns of CDT . $\{OC\}$ is a BPK-set and a BUC-set, and $\{OC\} \mapsto CT$ is a *known functional dependency*. If WITH ORDINALITY is not specified, then these rules do not identify any BPK-set, BUC-set, or non-axiomatic known functional dependency.

4.18.8 Known functional dependencies in a <table factor>

Let R be the result of <table factor> TF . Let S be the result of <table primary> immediately contained in TF . The counterparts of the BPK-sets and BUC-sets of S are the BPK-sets and BUCsets, respectively, of R . If $A \mapsto B$ is a functional dependency in S , and AC and BC are the counterparts of A and B , respectively, then $AC \mapsto BC$ is a *known functional dependency* in R .

4.18.9 Known functional dependencies in a <table reference>

Let R be the result of some <table reference> TR . The BPK-sets, BUC-sets, and functional dependencies of R are determined as follows.

Case:

- If TR immediately contains a <table factor> TF , then the counterparts of the BPK-sets and BUC-sets of TF are the BPK-sets and BUC-sets, respectively, of R . If $A \mapsto B$ is a functional dependency in the result of TF , and AC and BC are the counterparts of A and B , respectively, then $AC \mapsto BC$ is a *known functional dependency* in R .
- If TR immediately contains a <joined table> JT , then the counterparts of the BPK-sets and BUC-sets of JT are the BPK-sets and BUC-sets, respectively, of R . If $A \mapsto B$ is a functional dependency in the result of JT , and AC and BC are the counterparts of A and B , respectively, then $AC \mapsto BC$ is a *known functional dependency* in R .

4.18.10 Known functional dependencies in the result of a <from clause>

Let R be the result of some <from clause> FC .

If there is only one <table reference> TR in FC , then the counterparts of the BPK-sets of TR and the counterparts of the BUC-sets of TR are the BPK-sets and BUC-sets of TR , respectively. Otherwise, these rules do not identify any BPK-sets or BUC-sets in the result of FC .

If T is a <table reference> immediately contained in the <table reference list> of FC , then all known functional dependencies in T are *known functional dependencies* in R .

4.18.11 Known functional dependencies in the result of a <where clause>

Let T be the table that is the operand of the <where clause>. Let R be the result of the <where clause>. A set of columns S in R is a *BUC-set* if there is a <table reference> TR such that every member of S has a counterpart in TR , the counterpart of S in TR is a BUC-set, and $S \mapsto CR$, where CR is the set of all columns of R . If, in addition, the counterpart of S is a BPK-set, then S is a *BPK-set*.

If $A \mapsto B$ is a known functional dependency in T , then let AC be the set of columns of R whose counterparts are in A , and let BC be the set of columns of R whose counterparts are in B . $AC \mapsto BC$ is a *known functional dependency* in R .

If AP is an equality AND-component of the <search condition> simply contained in the <where clause> and one comparand of AP is a column reference CR , and the other comparand of AP is a <literal>, then let CRC be the counterpart of CR in R . $\{\} \mapsto \{CRC\}$ is a *known functional dependency* in R , where $\{\}$ denotes the empty set.

NOTE 50 — An SQL-implementation may also choose to recognize $\{\} \mapsto \{CRC\}$ as a known functional dependency if the other comparand is a deterministic expression containing no column references.

If AP is an equality AND-component of the <search condition> simply contained in the <where clause> and one comparand of AP is a column reference CRA , and the other comparand of AP is a column references CRB ,

then let $CRAC$ and $CRBC$ be the counterparts of CRA and CRB in R . $\{CRBC\} \mapsto \{CRAC\}$ and $\{CRAC\} \mapsto \{CRBC\}$ are *known functional dependencies* in R .

NOTE 51 — An SQL-implementation may also choose to recognize known functional dependencies of the form $\{CRAC_1, \dots, CRAC_N\} \mapsto \{CRBC\}$ if one comparand is a deterministic expressions that contains column references CRA_1, \dots, CRA_N and the other comparand is a column reference CRB .

4.18.12 Known functional dependencies in the result of a <group by clause>

Let $T1$ be the table that is the operand of the <group by clause>, and let R be the result of the <group by clause>.

Let G be the set of columns specified by the <grouping column reference list> of the <group by clause>, after applying all syntactic transformations to eliminate ROLLUP, CUBE, and GROUPING SETS.

The columns of R are the columns of G , with an additional column CI , whose value in any particular row of R somehow denotes the subset of rows of $T1$ that is associated with the combined value of the columns of G in that row.

If every element of G is a column reference to a known not null column, then G is a *BUC-set* of R . If G is a subset of a *BPK-set* of columns of $T1$, then G is a *BPK-set* of R .

$G \mapsto CI$ is a *known functional dependency* in R .

NOTE 52 — Any <set function specification> that is specified in conjunction with R is necessarily a function of CI . If $SFVC$ denotes the column containing the results of such a <set function specification>, then $CI \mapsto SFVC$ holds true, and it follows that $G \mapsto SFVC$ is a *known functional dependency* in the table containing $SFVC$.

4.18.13 Known functional dependencies in the result of a <having clause>

Let $T1$ be the table that is the operand of the <having clause>, let SC be the <search condition> simply contained in the <having clause>, and let R be the result of the <having clause>.

If S is a set of columns of R and the counterpart of S in $T1$ is a *BPK-set*, then S is a *BPK-set*. If the counterpart of S in $T1$ is a *BUC-set*, then S is a *BUC-set*.

Any known functional dependency in the <query expression>

```
SELECT * FROM T1 WHERE SC
```

is a *known functional dependency* in R .

4.18.14 Known functional dependencies in a <query specification>

Let T be the <table expression> simply contained in the <query specification> QS and let R be the result of the <query specification>.

Let SL be the <select list> of the <query specification>.

Let $T1$ be T extended to the right with columns arising from <value expression>s contained in the <select list>, as follows: A <value expression> VE that is not a column reference specifies a computed column CC in $T1$. For every row in $T1$, the value in CC is the result of VE .

Let S be a set of columns of R such that every element of S arises from the use of <asterisk> in SL or by the specification of a column reference as a <value expression> simply contained in SL . S has counterparts in T and $T1$. If the counterpart of S in T is a BPK-set, then S is a *BPK-set*. If the counterpart of S in T is a BUC-set or a BPK-set, then S is a *BUC-set*.

If $A \mapsto B$ is some known functional dependency in T , then $A \mapsto B$ is a *known functional dependency* in $T1$.

Let CC be the column specified by some <value expression> VE that is not possibly non-deterministic in the <select list>.

Let $OP1, OP2, \dots$ be the operands of VE that are column references whose qualifying query is QS and that are not contained in an aggregated argument of a <set function specification>.

If VE does not contain a <set function specification> whose aggregation query is QS , then $\{OP1, OP2, \dots\} \mapsto CC$ is a *known functional dependency* in $T1$.

If VE contains a <set function specification> whose aggregation query is QS , then let $\{G1, \dots\}$ be the set of grouping columns of T . $\{G1, \dots, OP1, OP2, \dots\} \mapsto CC$ is a *known functional dependency* in $T1$.

Let $C \mapsto D$ be some known functional dependency in $T1$. If all the columns of C have counterparts in R , then let DR be the set comprising those columns of D that have counterparts in R . $C \mapsto DR$ is a *known functional dependency* in R .

4.18.15 Known functional dependencies in a <query expression>

If a <with clause> is specified, and RECURSIVE is not specified, then the *BPK-sets*, *BUC-sets*, and *known functional dependencies* of the table identified by a <query name> in the <with list> are the same as the BPK-sets, BUC-sets, and known functional dependencies of the corresponding <query expression>, respectively. If RECURSIVE is specified, then the BPK-sets, BUC-sets, and non-axiomatic known functional dependencies are implementation-defined.

A <query expression> that is a <query term> that is a <query primary> that is a <simple table> is covered by previous Subclauses of this Clause.

If the <query expression> specifies UNION, EXCEPT or INTERSECT, then let $T1$ and $T2$ be the left and right operand tables and let R be the result. Let CR be the set comprising all the columns of R .

Each column of R has a counterpart in $T1$ and a counterpart in $T2$.

Case:

- If EXCEPT is specified, then a set S of columns of R is a *BPK-set* if its counterpart in $T1$ is a BPK-set. S is a *BUC-set* if its counterpart in $T1$ is a BUC-set.
- If UNION is specified, then there are no BPK-sets and no BUC-sets.
- If INTERSECT is specified, then a set S of columns of R is a *BPK-set* if either of its counterparts in $T1$ and $T2$ is a BPK-set. S is a *BUC-set* if either of its counterparts in $T1$ and $T2$ is a BUC-set.

Case:

4.18 Functional dependencies

- If UNION is specified, then no non-axiomatic functional dependency in $T1$ or $T2$ is a known functional dependency in R , apart from any functional dependencies determined by implementation-defined rules.
- If EXCEPT is specified, then all known functional dependencies in $T1$ are *known functional dependencies* in R .
- If INTERSECT is specified, then all known functional dependencies in $T1$ and all known functional dependencies in $T2$ are *known functional dependencies* in R .

NOTE 53 — Other known functional dependencies may be determined according to implementation-defined rules.

4.19 Candidate keys

If the functional dependency $CK \rightarrow CT$ holds true in some table T , where CT consists of all columns of T , and there is no proper subset $CK1$ of CK such that $CK1 \rightarrow CT$ holds true in T , then CK is a *candidate key* of T . The set of candidate keys SCK is nonempty because, if no proper subset of CT is a candidate key, then CT is a candidate key.

NOTE 54 — Because a candidate key is a set (of columns), SCK is therefore a set of sets (of columns).

A candidate key CK is a *strong candidate key* if CK is a BUC-set, or if T is a grouped table and CK is a subset of the set of grouping columns of T . Let $SSCK$ be the set of strong candidate keys.

Let PCK be the set of P such that P is a member of SCK and P is a *BPK-set*.

Case:

- If PCK is nonempty, then the *primary key* is chosen from PCK as follows: If PCK has exactly one element, then that element is the primary key; otherwise, the left-most element of PCK is chosen according to the “left-most rule” below. The primary key is also the *preferred candidate key*.
- Otherwise, there is no primary key and the *preferred candidate key* is chosen as follows.

Case:

- If $SSCK$ has exactly one element, then it is the preferred candidate key; otherwise, if $SSCK$ has more than one element, then the left-most element of $SSCK$ is chosen, according to the “left-most” rule below.
- Otherwise, if SCK has exactly one element, then it is the preferred candidate key; otherwise, the left-most element of SCK is chosen, according to the “left-most” rule below.
- The “left-most” rule:
 - This rule uses the ordering of the columns of a table, as specified elsewhere in this part of ISO/IEC 9075.

To determine the left-most of two sets of columns of T , first list each set in the order of the column-numbers of its members, extending the shorter list with zeros to the length of the longer list. Then, starting at the left of each ordered list, step forward until a pair of unequal column numbers, one from the same position in each list, is found. The list containing the number that is the smaller member of this pair identifies the left-most of the two sets of columns of T .

To determine the left-most of more than two sets of columns of T , take the left-most of any two sets, then pair that with one of the remaining sets and take the left-most, and so on until there are no remaining sets.

4.20 SQL-schemas

*This Subclause is modified by Subclause 4.3, “SQL-schemas”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 4.12, “SQL-schemas”, in ISO/IEC 9075-9.*

04 An SQL-schema is a persistent descriptor that includes:

- The name of the SQL-schema.
- The <authorization identifier> of the owner of the SQL-schema.
- The name of the default character set for the SQL-schema.
- The <schema path specification> defining the SQL-path for SQL-invoked routines for the SQL-schema.
- The descriptor of every component of the SQL-schema.

In this part of ISO/IEC 9075, the term “schema” is used only in the sense of SQL-schema. The persistent objects described by the descriptors are said to be *owned by* or to have been *created by* the <authorization identifier> of the schema. Each component descriptor is one of:

- A domain descriptor.
- A base table descriptor.
- A view descriptor.
- A constraint descriptor.
- A trigger descriptor.
- A privilege descriptor.
- A character set descriptor.
- A collation descriptor.
- A transliteration descriptor.
- A user-defined type descriptor.
- A routine descriptor.
- A sequence generator descriptor.

A schema is created initially using a <schema definition> and may be subsequently modified incrementally over time by the execution of <SQL schema statement>s. <schema name>s are unique within a catalog.

A <schema name> is explicitly or implicitly qualified by a <catalog name> that identifies a catalog.

09 Base tables and views are identified by <table name>s. A <table name> consists of a <schema name>, followed by a <period>, followed by an <identifier>. The <schema name> identifies the schema that includes the table descriptor of the base table or view identified by the <table name>. The <table name>s of base tables and views defined in different schemas can have equivalent <identifier>s.

NOTE 55 — Equivalence of <identifier>s is defined in Subclause 5.2, “<token> and <separator>”.

If a reference to a <table name> does not explicitly contain a <schema name>, then a specific <schema name> is implied. The particular <schema name> associated with such a <table name> depends on the context in which the <table name> appears and is governed by the rules for <schema qualified name>.

If a reference to an SQL-invoked routine that is contained in a <routine invocation> does not explicitly contain a <schema name>, then the SQL-invoked routine is selected from the SQL-path of the schema.

The *containing schema* of an <SQL schema statement> is defined as the schema identified by the <schema name> implicitly or explicitly contained in the name of the object that is created or manipulated by that SQL-statement.

4.21 Sequence generators

4.21.1 General description of sequence generators

A *sequence generator* is a mechanism for generating successive exact numeric values, one at a time. A sequence generator is either an *external sequence generator* or an *internal sequence generator*. An external sequence generator is a named schema object while an internal sequence generator is a component of another schema object. A sequence generator has a data type, which shall be an exact numeric type with scale 0 (zero), a minimum value, a maximum value, a start value, an increment, and a cycle option.

Specification of a sequence generator can optionally include the specification of a data type, a minimum value, a maximum value, a start value, an increment, and a cycle option.

If a sequence generator is associated with a negative increment, then it is a *descending sequence generator*; otherwise, it is an *ascending sequence generator*.

A sequence generator has a time-varying *current base value*, which is a value of its data type. A sequence generator has a cycle which consists of all the possible values between the minimum value and the maximum value which are expressible as $(\text{current base value} + N * \text{increment})$, where N is a non-negative number.

When a sequence generator is created, its current base value is initialized to the start value. Subsequently, the current base value is set to the value of the lowest non-issued value in the cycle for an ascending sequence generator, or the highest non-issued value in the cycle for a descending sequence generator.

Any time after a sequence generator is created, its current base value can be set to an arbitrary value of its data type by an <alter sequence generator statement>.

Changes to the current base value of a sequence generator are not controlled by SQL-transactions; therefore, commits and rollbacks of SQL-transactions have no effect on the current base value of a sequence generator.

A sequence generator is described by a sequence generator descriptor. A sequence generator descriptor includes:

- The sequence generator name that is a schema-qualified sequence generator name for an external sequence generator and a zero-length character string for an internal sequence generator.
- The data type descriptor of the data type associated with the sequence generator.
- The increment of the sequence generator.
- The maximum value of the sequence generator.

- The minimum value of the sequence generator.
- The cycle option of the sequence generator.
- The current base value of the sequence generator.

4.21.2 Operations involving sequence generators

When a <next value expression> is applied to a sequence generator *SG*, *SG* issues a value *V* taken from *SG*'s current cycle such that *V* is expressible as the current base value of *SG* plus *N* multiplied by the increment of *SG*, where *N* is a non-negative number.

Thus a sequence generator will normally issue all of the values in its cycle and these will normally be in increasing or decreasing order (depending on the sign of the increment) but within that general ordering separate subgroups of ordered values may occur.

If the sequence generator's cycle is exhausted (*i.e.*, it cannot issue a value that meets the criteria), then a new cycle is created with the current base value set to the minimum value of *SG* (if *SG* is an ascending sequence generator) or the maximum value of *SG* (if *SG* is a descending sequence generator).

If a new cycle is created and the descriptor of *SG* includes NO CYCLE, then an exception condition is raised.

If there are multiple instances of <next value expression>s specifying the same sequence generator within a single SQL-statement, all those instances return the same value for a given row processed by that SQL-statement.

4.22 SQL-client modules

An *SQL-client module* is an SQL-environment object that can include externally-invoked procedures and certain descriptors. An SQL-client module is created and destroyed by implementation-defined mechanisms (which can include the granting and revoking of privileges required for the use of the SQL-client module). An SQL-client module exists in the SQL-environment containing an SQL-client.

If an SQL-client module *S* is defined by an <SQL-client module definition> that contains a <module authorization identifier> *MAI*, then the owner of *S* is *MAI*; otherwise, *S* has no owner.

An SQL-client module can be specified by a <SQL-client module definition> (see [Subclause 13.1](#), “<SQL-client module definition>”).

An SQL-client module includes:

- The name, if any of the SQL-client module.
- The name of the host language from a compilation unit of which an externally-invoked procedure included in the module can be invoked.
- The <module authorization identifier>, if any.
- An indication of whether or not the <module authorization identifier> is to apply to execution of prepared statements resulting from invocation of externally-invoked procedures in the SQL-client module that contain <prepare statement>s or <execute immediate statement>s.

- SQL-client module defaults, for use in the application of Syntax Rules to <externally-invoked procedure>s, <temporary table declaration>s, and <declare cursor>s.
 - The name of the schema for use as the default <schema name> when deriving externally-invoked procedures from <externally-invoked procedure>s, specified either by the <schema name> or, failing that, by the <module authorization identifier>.
 - The SQL-path, if any, used to qualify:
 - Unqualified <routine name>s that are immediately contained in <routine invocation>s that are contained in the SQL-client module.
 - Unqualified <user-defined type name>s that are immediately contained in <path-resolved user-defined type name>s that are contained in the SQL-client module.
 - The names of zero or more *SQL-client module collations*, each specifying a collation for one or more character sets for the SQL-client module.
- The name, if specified, of the character set used to express the <SQL-client module definition>.

NOTE 56 — The <module character set specification> has no effect on the SQL language contained in the SQL-client module and exists only for compatibility with ISO/IEC 9075:1992. It may be used to document the character set of the SQL-client module.
- Module contents:
 - Zero or more temporary table descriptors.
 - Zero or more cursor declaration descriptors.
 - One or more externally-invoked procedures.

A compilation unit is a segment of executable code, possibly consisting of one or more subprograms. An SQL-client module is associated with a compilation unit during its execution. A single SQL-client module may be associated with multiple compilation units and multiple SQL-client modules may be associated with a single compilation unit. The manner in which this association is specified, including the possible requirement for execution of some implementation-defined statement, is implementation-defined. Whether a compilation unit may invoke or transfer control to other compilation units, written in the same or a different programming language, is implementation-defined.

4.23 Embedded syntax

This Subclause is modified by Subclause 4.1, “Embedded syntax”, in ISO/IEC 9075-10.

10 An <embedded SQL host program> is a compilation unit that consists of host language text and SQL text. The host language text conforms to the requirements of a specific programming language. The SQL text consists of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s, as defined in this International Standard. This allows database applications to be expressed in a hybrid form in which SQL-statements are embedded directly in a compilation unit. Such a hybrid compilation unit is defined to be equivalent to:

- An SQL-client module, containing externally-invoked procedures and declarations.

- A host language compilation unit in which each SQL-statement has been replaced by an invocation of an externally-invoked procedure in the SQL-client module, and the declarations contained in such SQL-statements have been suitably transformed into declarations in the host language.

If an <embedded SQL host program> contains an <embedded authorization declaration>, then it shall be the first statement or declaration in the <embedded SQL host program>. The <embedded authorization declaration> is not replaced by a procedure or subroutine call of an <externally-invoked procedure>, but is removed and replaced by syntax associated with the <SQL-client module definition>'s <module authorization clause>.

An implementation may reserve portions of the name spaces in the <embedded SQL host program> for the names of host language elements and host language environment elements that are generated to replace SQL-statements. This includes the names of procedures, subroutines, classes, and other resources in the host language environment, and for program variables and branch labels that may be generated as required to support the calling of these procedures or subroutines. Whether any such reservations are made, and the details of any such reservations, if any, is implementation-defined.

4.24 Dynamic SQL concepts

4.24.1 Overview of dynamic SQL

In many cases, the SQL-statement to be executed can be coded into an <SQL-client module definition> or into a compilation unit using the embedded syntax. In other cases, the SQL-statement is not known when the program is written and will be generated during program execution.

Dynamic execution of SQL-statements can generally be accomplished in two different ways. Statements can be *prepared* for execution and then later executed one or more times; when the statement is no longer needed for execution, it can be *released* by the use of a <deallocate prepared statement>. Alternatively, a statement that is needed only once can be executed without the preparation step—it can be *executed immediately* (not all SQL-statements can be executed immediately).

When a prepared statement is executed, if it has an owner, then it is executed under definer's rights; otherwise, it is executed under invoker's rights.

Many SQL-statements can be written to use “parameters” (which are manifested in static execution of SQL-statements as host parameters in <SQL procedure statement>s contained in <externally-invoked procedure>s in <SQL-client module definition>s or as host variables in <embedded SQL statement>s contained in <embedded SQL host program>s). In SQL-statements that are executed dynamically, the parameters are called dynamic parameters (<dynamic parameter specification>s) and are represented in SQL language by a <question mark> (?).

In many situations, an application that generates an SQL-statement for dynamic execution knows in detail the required characteristics (*e.g.*, <data type>, <length>, <precision>, <scale>, *etc.*) of each of the dynamic parameters used in the statement; similarly, the application may also know in detail the characteristics of the values that will be returned by execution of the statement. However, in other cases, the application may not know this information to the required level of detail; it is possible in some cases for the application to ascertain the information from the Information Schema, but in other cases (*e.g.*, when a returned value is derived from a computation instead of simply from a column in a table, or when dynamic parameters are supplied) this information is not generally available except in the context of preparing the statement for execution.

NOTE 57 — The Information Schema is defined in [ISO9075-11].

To provide the necessary information to applications, SQL permits an application to request the SQL-server to *describe* a prepared statement. The description of a statement identifies the number of input dynamic parameters (*describe input*) and their data type information or it identifies the number of output dynamic parameters or values to be returned (*describe output*) and their data type information. The description of a statement is placed into the SQL descriptor areas already mentioned.

Many, but not all, SQL-statements can be prepared and executed dynamically.

NOTE 58 — The complete list of statements that may be dynamically prepared and executed is defined in [Subclause 4.33.7, “Preparable and immediately executable SQL-statements”](#).

Certain “set statements” (<set catalog statement>, <set schema statement>, <set names statement>, and <set path statement>) have no effect other than to set up default information (catalog name, schema name, character set, and SQL path, respectively) to be applied to other SQL-statements that are prepared or executed immediately or that are invoked directly.

Syntax errors and Access Rule violations caused by the preparation or immediate execution of <preparable statement>s are identified when the statement is prepared (by <prepare statement>) or when it is executed (by <execute statement> or <execute immediate statement>); such violations are indicated by the raising of an exception condition.

4.24.2 Dynamic SQL statements and descriptor areas

An <execute immediate statement> can be used for a one-time preparation and execution of an SQL-statement. A <prepare statement> is used to prepare the generated SQL-statement for subsequent execution. A <deallocate prepared statement> is used to deallocate SQL-statements that have been prepared with a <prepare statement>. A description of the input dynamic parameters for a prepared statement can be obtained by execution of a <describe input statement>. A description of the resultant columns of a <dynamic select statement> or <dynamic single row select statement> can be obtained by execution of a <describe output statement>. A description of the output dynamic parameters of a statement that is neither a <dynamic select statement> nor a <dynamic single row select statement> can be obtained by execution of a <describe output statement>.

For a statement other than a <dynamic select statement>, an <execute statement> is used to associate parameters with the prepared statement and execute it as though it had been coded when the program was written. For a <dynamic select statement>, the prepared <cursor specification> is associated with a declared dynamic cursor via a <dynamic declare cursor> or with an extended statement cursor via an <allocate cursor statement>. The dynamic cursor can be opened and dynamic parameters can be associated with the dynamic cursor with a <dynamic open statement>. Operations on an open dynamic cursor are described in [Subclause 4.32.2, “Operations on and using cursors”](#).

The interface for input dynamic parameters and output dynamic parameters for a prepared statement and for the resulting values from a <dynamic fetch statement> or the execution of a prepared <dynamic single row select statement> can be either a list of dynamic parameters or embedded variables or an SQL descriptor area. An SQL descriptor area consists of one or more item descriptor areas, together with a header that includes a count of the number of those item descriptor areas. The header of an SQL descriptor area consists of the fields in [Table 23, “Data types of <key word>s used in the header of SQL descriptor areas”](#), in [Subclause 20.1, “Description of SQL descriptor areas”](#). Each item descriptor area consists of the fields specified in [Table 24, “Data types of <key word>s used in SQL item descriptor areas”](#), in [Subclause 20.1, “Description of SQL descriptor areas”](#). The SQL descriptor area is allocated and maintained by the system with the following statements: <allocate descriptor statement>, <deallocate descriptor statement>, <set descriptor statement>, and <get descriptor statement>.

Two kinds of identifier are used for referencing dynamic SQL objects, *extended names* and *non-extended names*. An extended name is an <identifier> assigned to a parameter or variable and the object it identifies is referenced indirectly, by referencing that parameter or variable. A non-extended name is just an <identifier> and the object it identifies is referenced by using that <identifier> directly in an SQL-statement.

SQL descriptor areas, dynamic statements, and dynamic cursors can be identified either by non-extended names or by extended names.

Two extended names are equivalent if their values, with leading and trailing <space>s removed, are equivalent according to the rules for <identifier> comparison in Subclause 5.2, “<token> and <separator>”.

The *scope* of an extended name is either *global* or *local* and is determined by the run-time context in which the object it identifies is brought into existence.

The scope of a global extended name *GEN* is the SQL-session, meaning that, during the existence of the object *O* it identifies, *GEN* can be used to reference *O* by any SQL-statement executed in that SQL-session.

The scope of a local extended name *LEN* is the SQL-client module *M* containing the externally-invoked procedure that is being executed when the object *O* identified by *LEN* is brought into existence. This means that, during the existence of *O*, *LEN* can be used to reference *O* by any SQL-statement executed in the same SQL-session by an externally-invoked procedure in *M*.

The scope of a non-extended name is the <SQL-client module definition> containing the SQL-statement that defines it.

NOTE 59 — The namespace of non-extended names is different from the namespace of extended names.

Let *PRP* be the prepared statement resulting from execution of a <prepare statement> in an externally-invoked procedure, SQL-invoked routine, or triggered action *E*. In the following cases, *PRP* has no owner:

- *E* is an SQL-invoked routine whose security characteristic is INVOKED.
- *E* is an externally-invoked procedure contained in an SQL-client module that either has no owner or for which FOR STATIC ONLY was specified.

Otherwise, the owner of *PRP* is the owner of *E*.

4.25 Direct invocation of SQL

Direct invocation of SQL is a mechanism for executing direct SQL-statements, known as <direct SQL statement>s. In direct invocation of SQL, the method of invoking <direct SQL statement>s, the method of raising conditions that result from the execution of <direct SQL statement>s, the method of accessing the diagnostics information that results from the execution of <direct SQL statement>s, and the method of returning the results are implementation-defined.

4.26 Externally-invoked procedures

An externally-invoked procedure consists of an SQL-statement and can be invoked from a compilation unit of a host language. The host language is specified by the <language clause> of the SQL-client module that contains the externally-invoked procedure.

4.27 SQL-invoked routines

4.27.1 Overview of SQL-invoked routines

*This Subclause is modified by Subclause 4.4.1, “Overview of SQL-invoked routines”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 4.2.1, “Overview of SQL-invoked routines”, in ISO/IEC 9075-13.*

An *SQL-invoked routine* is an SQL-invoked procedure or an SQL-invoked function. An SQL-invoked routine comprises at least a <schema qualified routine name>, a sequence of <SQL parameter declaration>s, and a <routine body>.

04 An SQL-invoked routine is an element of an SQL-schema and is called a *schema-level routine*.

An SQL-invoked routine *SR* is said to be *dependent* on a user-defined type *UDT* if *SR* is created during the execution of the <user-defined type definition> that created *UDT* or if *SR* is created during the execution of an <alter type statement> that specifies an <add attribute definition>. An SQL-invoked routine that is dependent on a user-defined type cannot be modified by an <alter routine statement> or be destroyed by a <drop routine statement>. It is destroyed implicitly by a <drop data type statement>.

An *SQL-invoked procedure* is an SQL-invoked routine that is invoked from an SQL <call statement>. An SQL-invoked procedure may have input SQL parameters, output SQL parameters, and SQL parameters that are both input SQL parameters and output SQL parameters. The format of an SQL-invoked procedure is specified by <SQL-invoked procedure> (see Subclause 11.51, “<SQL-invoked routine>”).

An SQL-invoked procedure may optionally be specified to require a new savepoint level to be established when it is invoked and destroyed on return from the executed routine body. The alternative of not taking a savepoint can also be directly specified with OLD SAVEPOINT LEVEL. When an SQL-invoked function is invoked a new savepoint level is always established. Savepoint levels are described in Subclause 4.35.2, “Savepoints”.

An *SQL-invoked function* is an SQL-invoked routine whose invocation returns a value. Every parameter of an SQL-invoked function is an input SQL parameter, one of which may be designated as the result SQL parameter. The format of an SQL-invoked function is specified by <SQL-invoked function> (see Subclause 11.51, “<SQL-invoked routine>”). An SQL-invoked function can be a *type-preserving function*; a type-preserving function is an SQL-invoked function that has a result SQL parameter. The most specific type of a non-null result of invoking a type-preserving function shall be compatible with the most specific type of the value of the argument substituted for its result SQL parameter.

An *SQL-invoked method* is an SQL-invoked function that is specified by <method specification designator> (see Subclause 11.51, “<SQL-invoked routine>”). There are three kinds of SQL-invoked methods: *SQL-invoked constructor methods*, *instance SQL-invoked methods* and *static SQL-invoked methods*. All SQL-invoked methods are associated with a user-defined type, also known as the *type of the method*. The <method characteristic>s of an SQL-invoked method are specified by a <method specification> contained in the <user-defined type definition> of the type of the method. Both an instance SQL-invoked method and an SQL-invoked constructor method satisfy the following conditions:

- Its first parameter, called the *subject parameter*, has a declared type that is a user-defined type. The type of the subject parameter is the type of the method. A parameter other than the subject parameter is called an *additional parameter*.
- Its descriptor is in the same schema as the descriptor of the data type of its subject parameter.

An SQL-invoked constructor method satisfies the following additional conditions:

- Its <method name> is equivalent to the <qualified identifier> simply contained in the <user-defined type name> included in the user-defined type descriptor of the type of the method.

13 A static SQL-invoked method satisfies the following conditions:

- It has no subject parameter. Its first parameter, if any, is treated no differently than any other parameter.
- Its descriptor is in the same schema as the descriptor of the structured type of the method. The name of this type (or of some subtype of it) is always specified together with the name of the method when the method is to be invoked.

An SQL-invoked function that is not an SQL-invoked method is an *SQL-invoked regular function*. An SQL-invoked regular function is specified by <function specification> (see [Subclause 11.51](#), “<SQL-invoked routine>”).

A *null-call function* is an SQL-invoked function that is defined to return the null value if any of its input arguments is the null value. A null-call function is an SQL-invoked function whose <null-call clause> specifies “RETURNS NULL ON NULL INPUT”.

4.27.2 Characteristics of SQL-invoked routines

This Subclause is modified by Subclause 4.2.2, “Characteristics of SQL-invoked routines”, in ISO/IEC 9075-13.

An SQL-invoked routine can be an *SQL routine* or an *external routine*. An SQL routine is an SQL-invoked routine whose <language clause> specifies SQL. The <routine body> of an SQL routine is an <SQL procedure statement>; the <SQL procedure statement> forming the <routine body> can be any SQL-statement, including an <SQL control statement>, but excluding an <SQL connection statement> and an <SQL transaction statement> other than a <savepoint statement>, a <release savepoint statement>, or a <rollback statement> that specifies a <savepoint clause>.

13 An external routine is one whose <language clause> does not specify SQL. The <routine body> of an external routine is an <external body reference> whose <external routine name> identifies a program written in some programming language other than SQL. The program identified by <external routine name> shall not execute either an <SQL connection statement> or an <SQL transaction statement> other than a <savepoint statement>, a <release savepoint statement>, or a <rollback statement> that specifies a <savepoint clause>.

An SQL-invoked routine is uniquely identified by a <specific name>, called the *specific name* of the SQL-invoked routine.

13 SQL-invoked routines are invoked differently depending on their form. SQL-invoked procedures are invoked by <call statement>s. SQL-invoked regular functions are invoked by <routine invocation>s. Instance SQL-invoked methods are invoked by <method invocation>s, while SQL-invoked constructor methods are invoked by <new specification>s and static SQL-invoked methods are invoked by <static method invocation>s. An invocation of an SQL-invoked routine specifies the <routine name> of the SQL-invoked routine and supplies a sequence of argument values corresponding to the <SQL parameter declaration>s of the SQL-invoked routine. A *subject routine* of an invocation is an SQL-invoked routine that may be invoked by a <routine invocation>. After the selection of the subject routine of a <routine invocation>, the SQL arguments are evaluated and the SQL-invoked routine that will be executed is selected. If the subject routine is an instance SQL-invoked method, then the SQL-invoked routine that is executed is selected from the set of overriding methods of the subject routine. (The term “set of overriding methods” is defined in the General Rules of [Subclause 10.4](#), “<routine invocation>”). The overriding method that is selected is the overriding method with a subject parameter the type designator of whose declared type precedes that of the declared type of the subject parameter of every

other overriding method in the type precedence list of the most specific type of the value of the SQL argument that corresponds to the subject parameter. See the General Rules of [Subclause 10.4, “<routine invocation>”](#). If the subject routine is not an SQL-invoked method, then the SQL-invoked routine executed is that subject routine. After the selection of the SQL-invoked routine for execution, the values of the SQL arguments are assigned to the corresponding SQL parameters of the SQL-invoked routine and its <routine body> is executed. If the SQL-invoked routine is an SQL routine, then the <routine body> is an <SQL procedure statement> that is executed according to the General Rules of <SQL procedure statement>. If the SQL-invoked routine is an external routine, then the <routine body> identifies a program written in some programming language other than SQL that is executed according to the rules of that host language.

The <routine body> of an SQL-invoked routine is always executed under the same SQL-session from which the SQL-invoked routine was invoked. Before the execution of the <routine body>, a new context for the current SQL-session is created and the values of the current context preserved. When the execution of the <routine body> completes the original context of the current SQL-session is restored.

13 If the SQL-invoked routine is an external routine, then an effective SQL parameter list is constructed before the execution of the <routine body>. The effective SQL parameter list has different entries depending on the parameter passing style of the SQL-invoked routine. The value of each entry in the effective SQL parameter list is set according to the General Rules of [Subclause 10.4, “<routine invocation>”](#), and passed to the program identified by the <routine body> according to the rules of [Subclause 13.6, “Data type correspondences”](#). After the execution of that program, if the parameter passing style of the SQL-invoked routine is SQL, then the SQL-implementation obtains the values for output parameters (if any), the value (if any) returned from the program, the value of the SQLSTATE, and the value of the message text (if any) from the values assigned by the program to the effective SQL parameter list. If the parameter passing style of the SQL-invoked routine is GENERAL, then such values are obtained in an implementation-defined manner.

Different SQL-invoked routines can have equivalent <routine name>s. No two SQL-invoked functions in the same schema are allowed to have the same signature. No two SQL-invoked procedures in the same schema are allowed to have the same name and the same number of parameters. Subject routine determination is the process for choosing the subject routine for a given <routine invocation> given a <routine name> and an <SQL argument list>. Subject routine determination for SQL-invoked functions considers the most specific types of all of the arguments (that is, all of the arguments that are not <dynamic parameter specification>s whose types are not known at the time of subject routine determination) to the invocation of the SQL-invoked function in order from left to right. Where there is not an exact match between the most specific types of the arguments and the declared types of the parameters, type precedence lists are used to determine the closest match. See [Subclause 9.4, “Subject routine determination”](#).

If a <routine invocation> is contained in a <query expression> of a view, in the <search condition> of a check constraint or an assertion, the <triggered action> of a trigger, or in an <SQL-invoked routine>, then the subject routine for that invocation is determined at the time the view is created, the check constraint is defined, the assertion is created, the trigger is created, or the SQL-invoked routine is created. If the subject routine is an SQL-invoked procedure, an SQL-invoked regular function, or a static SQL-invoked method, then the same SQL-invoked routine is executed whenever the view is used, the check constraint or assertion is evaluated, the trigger is executed, or the SQL-invoked routine is invoked. If the subject routine is an instance SQL-invoked method, then the SQL-invoked routine that is executed is determined whenever the view is used, the check constraint or assertion is evaluated, the trigger is executed, or the SQL-invoked routine is invoked, based on the most specific type of the value resulting from the evaluation of the SQL argument that correspond to the subject parameter. See the General Rules of [Subclause 10.4, “<routine invocation>”](#).

All <identifier chain>s in the <routine body> of an SQL routine are resolved to identify the basis and basis referent at the time that the SQL routine is created. Thus, the same columns and SQL parameters are referenced whenever the SQL routine is invoked.

An SQL-invoked routine is either *deterministic* or *possibly non-deterministic*. An SQL-invoked function that is deterministic always returns the identical return value for a given list of SQL argument values. An SQL-invoked procedure that is deterministic always returns the identical values in its output and inout SQL parameters for a given list of SQL argument values. An SQL-invoked routine is possibly non-deterministic if it might produce nonidentical results when invoked with the identical list of SQL argument values.

An external routine *does not possibly contain SQL*, *possibly contains SQL*, *possibly reads SQL-data*, or *possibly modifies SQL-data*. Only an external routine that possibly contains SQL, *possibly reads SQL-data*, or *possibly modifies SQL-data* is permitted to execute SQL-statements during its invocation. Only an SQL-invoked routine that possibly reads SQL-data or possibly modifies SQL-data may read SQL-data during its invocation. Only an SQL-invoked routine that possibly modifies SQL-data may modify SQL-data during its invocation.

An SQL-invoked routine has a *routine authorization identifier*, which is (directly or indirectly) the authorization identifier of the owner of the schema that contains the SQL-invoked routine at the time that the SQL-invoked routine is created.

4.27.3 Execution of SQL-invoked routines

This Subclause is modified by Subclause 4.4.2, “Execution of SQL-invoked routines”, in ISO/IEC 9075-4.

When an SQL-invoked routine is invoked, a copy of the current SQL-session context is pushed onto the stack and some values are modified (see the General Rules of Subclause 10.4, “<routine invocation>”, before the <routine body> is executed. The treatment of the authorization stack is described in Subclause 4.34.1.1, “SQL-session authorization identifiers”.

04 An SQL-invoked routine has a *routine SQL-path*, which is inherited from its containing SQL-schema, the current SQL-session, or the containing SQL-client module.

An SQL-invoked routine that is an external routine also has an *external routine SQL-path*, which is derived from the <module path specification>, if any, of the <SQL-client module definition> contained in the external program identified by the routine body of the external routine. If that <SQL-client module definition> does not specify a <module path specification>, then the external routine SQL-path is an implementation-defined SQL-path. For both SQL and external routines, the SQL-path of the current SQL-session is used to determine the search order for the subject routine of a <routine invocation> whose <routine name> does not contain a <schema name> if the <routine invocation> is contained in a <preparable statement> that is prepared in the current SQL-session or in a <direct SQL statement>. SQL routines use the routine SQL-path to determine the search order for the subject routines of a <routine invocation> whose <routine name> does not contain a <schema name> if the <routine invocation> is not contained in a <preparable statement> that is prepared in the current SQL-session or in a <direct SQL statement>. External routines use the external routine SQL-path to determine the search order for the subject routine of a <routine invocation> whose <routine name> does not contain a <schema name> if the <routine invocation> is not contained in a <preparable statement> that is prepared in the current SQL-session or in a <direct SQL statement>.

4.27.4 Routine descriptors

This Subclause is modified by Subclause 4.4.3, “Routine descriptors”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 4.5.1, “Routine descriptors”, in ISO/IEC 9075-14.

04 An SQL-invoked routine is described by a *routine descriptor*. A routine descriptor includes:

- The routine name of the SQL-invoked routine.
- The specific name of the SQL-invoked routine.
- The routine authorization identifier of the SQL-invoked routine.
- The routine SQL-path of the SQL-invoked routine.
- The name of the language in which the body of the SQL-invoked routine is written.
- For each of the SQL-invoked routine's SQL parameters, the <SQL parameter name>, if it is specified, the <data type>, the ordinal position, and an indication of whether the SQL parameter is an input SQL parameter, an output SQL parameter, or both an input SQL parameter and an output SQL parameter.
- An indication of whether the SQL-invoked routine is an SQL-invoked function or an SQL-invoked procedure.
- If the SQL-invoked routine is an SQL-invoked procedure, then the maximum number of returned result sets.
- An indication of whether the SQL-invoked routine is deterministic or possibly non-deterministic.
- Indications of whether the SQL-invoked routine possibly modifies SQL-data, possibly reads SQL-data, possibly contains SQL, or does not possibly contain SQL.
- If the SQL-invoked routine is an SQL-invoked function, then:
 - The <returns data type> of the SQL-invoked function.
 - If the <returns data type> simply contains <locator indication>, then an indication that the return value is a locator.
 - An indication of whether the SQL-invoked function is a type-preserving function or not.
 - An indication of whether the SQL-invoked function is a mutator function or not.
 - If the SQL-invoked function is a type-preserving function, then an indication of which parameter is the result parameter.
 - An indication of whether the SQL-invoked function is a null-call function.
 - An indication of whether the SQL-invoked function is an SQL-invoked method.
- The creation timestamp.
- The last-altered timestamp.
- If the SQL-invoked routine is an SQL routine, then:
 - The SQL routine body of the SQL-invoked routine.
 - The SQL security characteristic of the SQL routine.
- If the SQL-invoked routine is an external routine, then:
 - The external routine name of the external routine.
 - The <parameter style> of the external routine.

- If the external routine specifies a <result cast>, then an indication that it specifies a <result cast> and the <data type> specified in the <result cast>. If <result cast> contains <locator indication>, then an indication that the <data type> specified in the <result cast> has a locator indication.
 - The external security characteristic of the external routine.
 - The external routine SQL-path of the external routine.
 - The effective SQL parameter list of the external routine.
 - For every SQL parameter that has an associated from-sql function *FSF*, the specific name of *FSF*.
 - For every SQL parameter that has an associated to-sql function *TSF*, the specific name of *TSF*.
 - If the SQL-invoked routine is an external function and if it has a to-sql function *TRF* associated with the result, then the specific name of *TRF*.
 - For every SQL parameter whose <SQL parameter declaration> contains <locator indication>, an indication that the SQL parameter is a locator parameter.
- The schema name of the schema that includes the SQL-invoked routine.
- If the SQL-invoked routine is an SQL-invoked method, then:
- The name of the user-defined type whose descriptor contains the corresponding method specification descriptor.
 - An indication of whether *STATIC* was specified.
- An indication of whether the SQL-invoked routine is dependent on a user-defined type.
- An indication as to whether or not the SQL-invoked routine requires a new savepoint level to be established when it is invoked.

4.27.5 Result sets returned by SQL-invoked procedures

This Subclause is modified by Subclause 4.5.1, “Result sets returned by SQL-invoked procedures”, in ISO/IEC 9075-3.

NOTE 60 — Subclause 3.1.6, “Definitions provided in Part 2”, gives definitions of the terms *result set*, *result set sequence*, *returned result set*, and *with-return cursor*, which are used in this Subclause.

An invocation of an SQL-invoked procedure *SIP1* might bring into existence a result set sequence *RSS*. *RSS* consists of the result sets of with-return cursors opened by *SIP1* and remaining open when *SIP1* terminates, placed in the order in which those result sets are created during the execution of *SIP1*.

NOTE 61 — If the same cursor is opened more than once during the execution of *SIP1*, then it is the last opening that is considered to create the result set, even if the result set in question is identical to that created by some earlier opening.

RSS is available through a received cursor, by executing a particular form of <allocate cursor statement>, only to the invoker *INV* of *SIP1*. If *SIP1* is invoked by an SQL-invoked procedure *SIP2*, then *INV* is *SIP2*. If *SIP1* is invoked by an externally-invoked procedure *EIP*, then *INV* is the SQL-client module containing *EIP*. Otherwise, *RSS* is not available.

NOTE 62 — Only the immediate invoker is considered. For example, if an externally-invoked procedure *EIP* executes a <call statement> invoking an SQL-invoked procedure *SIP3* that invokes *SIP1*, then the result set sequence returned by *SIP1* is available only to *SIP3*, until either *SIP3* returns control to *EIP* or another invocation of *SIP1* by *SIP3* is given before *SIP3* returns. There is

no mechanism whereby *SIP3* can return *SIP1*'s result set sequence to the invoker of *SIP3*, even if *SIP3* is defined to be able to return a result set sequence.

The invocation of *SIP1* by *INV* destroys any existing result set sequence that might have arisen from some previous invocation of *SIP1* by *INV*. All result set sequences available to *INV* are destroyed when *INV* terminates.

A returned result set includes a sequence of rows called the *constituent rows* and an *initial cursor position*. All rows of a result set sequence are of the same row type, called the *row type* of the result set. The *degree* of a result set is the number of columns in the result set's row type. The number of rows in the result set is its *cardinality*.

The constituent rows and initial cursor position of each returned result set *RS* in *RSS* are determined when *SIP1* returns to *INV*. If the with-return cursor *C* for *RS* is scrollable, then the constituent rows of *RS* are those of the result set of *C* as it exists when *SIP1* terminates; otherwise, the constituent rows are as for scrollable cursors, except that rows preceding the current cursor position of *C* are excluded. If *C* is scrollable, then the initial cursor position of *RS* is the position of *C* when *SIP1* terminates; otherwise, the initial cursor position of *RS* is before the first row.

NOTE 63 — The result set of *C* as it exists when *SIP1* terminates might differ from that generated when *C* was opened, if, for example, any <delete statement: positioned>s or <update statement: positioned>s are executed by *SIP1* before it terminates.

The maximum number of returned result sets that may form a result set sequence is specified by the <returned result sets characteristic> contained in the <SQL-invoked routine> defining *SIP1*. If the actual number of with-return cursors that remain open when *SIP1* returns is greater than the maximum number of returned result sets specified in the <returned result sets characteristic> clause, then a warning condition is raised. It is implementation-dependent whether or not result sets whose positions are greater than that maximum number are returned.

[03] A result set is described by a result set descriptor. A *result set descriptor* includes:

- A <cursor specification>.
- A sequence of rows.
- A position within the sequence of rows (before a row, on a row, or after the last row).
- The operational properties:
 - The operational sensitivity property (either SENSITIVE, INSENSITIVE, or ASENSITIVE).
 - The operational scrollability property (either SCROLL or NO SCROLL).
 - The operational holdability property (either WITH HOLD or WITHOUT HOLD).
 - The operational returnability property (either WITH RETURN or WITHOUT RETURN).
- If the operational holdability property is WITH HOLD and the position is on a row, then an indication whether a <fetch statement> or <dynamic fetch statement> has been issued since the last <commit statement>.

4.28 SQL-paths

This Subclause is modified by Subclause 4.5, “SQL-paths”, in ISO/IEC 9075-4.

An SQL-path is a list of one or more <schema name>s that determines the search order for one of the following:

- The subject routine of a <routine invocation> whose <routine name> does not contain a <schema name>.

— The user-defined type when the <path-resolved user-defined type name> does not contain a <schema name>.

[04] The value specified by CURRENT_PATH is the value of the SQL-path of the current SQL-session. This SQL-path is used to search for the subject routine of a <routine invocation> whose <routine name> does not contain a <schema name> when the <routine invocation> is contained in <preparable statement>s that are prepared in the current SQL-session by either an <execute immediate statement> or a <prepare statement>, or contained in <direct SQL statement>s that are invoked directly. The definition of SQL-schemas specifies an SQL-path that is used to search for the subject routine of a <routine invocation> whose <routine name>s do not contain a <schema name> when the <routine invocation> is contained in the <schema definition>.

4.29 Host parameters

4.29.1 Overview of host parameters

A host parameter is declared in an <externally-invoked procedure> by a <host parameter declaration>. A host parameter either assumes or supplies the value of the corresponding argument in the invocation of the <externally-invoked procedure>.

A <host parameter declaration> specifies the <data type> of its value, which maps to the host language type of its corresponding argument. Host parameters cannot be null, except through the use of indicator parameters.

4.29.2 Status parameters

This Subclause is modified by Subclause 4.6.1, “Status parameters”, in ISO/IEC 9075-4.

The SQLSTATE host parameter is a status parameter. It is set to status codes that indicate either that a call of the <externally-invoked procedure> completed successfully or that an exception condition was raised during execution of the <externally-invoked procedure>.

An <externally-invoked procedure> shall specify the SQLSTATE host parameter. The SQLSTATE host parameter is a character string host parameter for which exception values are defined in [Clause 24, “Status codes”](#).

If a condition is raised that causes a statement to have no effect other than that associated with raising the condition (that is, not a completion condition), then the condition is said to be an *exception condition* or *exception*. If a condition is raised that permits a statement to have an effect other than that associated with raising the condition (corresponding to an SQLSTATE class value of *successful completion*, *warning*, or *no data*), then the condition is said to be a *completion condition*.

Exception conditions or completion conditions may be raised during the execution of an <SQL procedure statement>. One of the conditions becomes the active condition when the <SQL procedure statement> terminates. If the active condition is an exception condition, then it will be called the active exception condition. If the active condition is a completion condition, then it will be called the active completion condition.

The completion condition *warning* is broadly defined as completion in which the effects are correct, but there is reason to caution the user about those effects. It is raised for implementation-defined conditions as well as

4.29 Host parameters

conditions specified in this part of ISO/IEC 9075. The completion condition *no data* has special significance and is used to indicate an empty result. The completion condition *successful completion* is defined to indicate a completion condition that does not correspond to *warning* or *no data*. This includes conditions in which the SQLSTATE subclass provides implementation-defined information of a non-cautionary nature.

For the purpose of choosing status parameter values to be returned, exception conditions for transaction rollback have precedence over exception conditions for statement failure. Similarly, the completion condition *no data* has precedence over the completion condition *warning*, which has precedence over the completion condition *successful completion*. All exception conditions have precedence over all completion conditions. The values assigned to SQLSTATE shall obey these precedence requirements.

4.29.3 Data parameters

A data parameter is a host parameter that is used to either assume or supply the value of data exchanged between a host program and an SQL-implementation.

4.29.4 Indicator parameters

An indicator parameter is an integer host parameter that is specified immediately following another host parameter. Its primary use is to indicate whether the value that the other host parameter assumes or supplies is a null value. An indicator host parameter cannot immediately follow another indicator host parameter.

The other use for indicator parameters is to indicate whether string data truncation occurred during a transfer between a host program and an SQL-implementation in host parameters or host variables. If a non-null string value is transferred and the length of the target is sufficient to accept the entire source value, then the indicator parameter or variable is set to 0 (zero) to indicate that truncation did not occur. However, if the length of the target is insufficient, the indicator parameter or variable is set to the length (in characters or octets, as appropriate) of the source value to indicate that truncation occurred and to indicate original length in characters or octets, as appropriate, of the source.

4.29.5 Locators

A host parameter, a host variable, an SQL parameter of an external routine, or the value returned by an external function may be specified to be a *locator* by specifying AS LOCATOR. A locator is an SQL-session object, rather than SQL-data, that can be used to reference an SQL-data instance. A locator is either a large object locator, a user-defined type locator, an array locator, or a multiset locator.

A large object locator is one of the following:

- Binary large object locator, a value of which identifies a binary large object.
- Character large object locator, a value of which identifies a large object character string.
- National character large object locator, a value of which identifies a national large object character string.

A user-defined type locator identifies a value of the user-defined type specified by the locator specification. An array locator identifies a value of the array type specified by the locator specification. A multiset locator identifies a value of the multiset type specified by the locator specification.

When the value at a site of binary large object type, character large object type, user-defined type, array type, or multiset type is to be assigned to a locator of the corresponding type, an implementation-dependent four-octet integer value is generated and assigned to the target. A locator value uniquely identifies a value of the corresponding type.

A locator may be either *valid* or *invalid*. A host parameter or host variable specified as a locator may be further specified to be a *holdable locator*. When a locator is initially created, it is marked valid and, if applicable, not holdable. A *<hold locator statement>* identifying the locator shall be specifically executed before the end of the SQL-transaction in which it was created in order to make that locator holdable.

A non-holdable locator remains valid until the end of the SQL-transaction in which it was generated, unless it is explicitly made invalid. A non-holdable locator becomes invalid whenever a *<free locator statement>* is executed or whenever a *<rollback statement>* that specifies a *<savepoint clause>* identifying a savepoint established prior to the generation of the locator is executed. Accessing an invalid locator causes an exception condition to be raised.

A holdable locator may remain valid beyond the end of the SQL-transaction in which it is generated. A holdable locator becomes invalid whenever a *<free locator statement>* identifying that locator is executed or the SQL-transaction in which it is generated or any subsequent SQL-transaction is rolled back. All locators remaining valid at the end of an SQL-session are marked invalid when that SQL-session terminates.

4.30 Diagnostics area

This Subclause is modified by Subclause 4.7, “Diagnostics area”, in ISO/IEC 9075-4.

A diagnostics area is a place where completion and exception condition information is stored when an SQL-statement is executed. The diagnostics areas associated with an SQL-session form the *diagnostics area stack* of that SQL-session. For definitional purposes, the diagnostics areas in this stack are considered to be numbered sequentially beginning with 1 (one). An additional diagnostics area is maintained by the SQL-client, as described in Subclause 4.2.3.1, “SQL-clients”, in [ISO9075-1].

Two operations on diagnostics area stacks are specified in this International Standard for definitional purposes only. *Pushing* a diagnostics area stack effectively creates a new first diagnostics area, incrementing the ordinal position of every existing diagnostics area in the stack by 1 (one). The content of the new first diagnostics area is initially a copy of the content of the old (now second) one. *Popping* a diagnostics area stack effectively destroys the first diagnostics area in the stack and decrements the ordinal position of every remaining diagnostics area by 1 (one). The maximum number of diagnostics areas in a diagnostics area stack is implementation-dependent.

Each diagnostics area consists of a *statement area* and a sequence of one or more *condition areas*, each of which is at any particular time either *occupied* or *vacant*. A diagnostics area is *empty* when each of its condition areas is vacant; *emptying* a diagnostics area brings about this state. A statement area consists of a collection of named *statement information items*. A condition area consists of a collection of named *condition information items*.

A statement information item gives information about the innermost SQL-statement that is being executed when a condition is raised. A condition information item gives information about the condition itself. The names and data types of statement and condition information items are given in Table 30, “<statement information

item name>s for use with <get diagnostics statement>”, and in Table 31, “<condition information item name>s for use with <get diagnostics statement>”. Their meanings are given by the General Rules of Subclause 23.1, “<get diagnostics statement>”.

At the beginning of the execution of any <SQL procedure statement> that is not an <SQL diagnostics statement>, the first diagnostics area is emptied. An implementation places information about a completion condition or an exception condition reported by SQLSTATE into a vacant condition area in this diagnostics area. If other conditions are raised, the extent to which these cause further condition areas to become occupied is implementation-defined.

An <externally-invoked procedure> containing an <SQL diagnostics statement> returns a code indicating a completion or an exception condition for that statement via SQLSTATE, but does not necessarily cause any vacant condition areas to become occupied.

The number of condition areas per diagnostics area is referred to as the *condition area limit*. An SQL-agent may set the condition area limit with the <set transaction statement>; if the SQL-agent does not specify the condition area limit, then the condition area limit is implementation-dependent, but shall be at least one condition area. An SQL-implementation may place information into this area about fewer conditions than there are condition areas. The ordering of the information about conditions placed into a diagnostics area is implementation-dependent, except that the first condition area in a diagnostics area always corresponds to the condition specified by the SQLSTATE value.

The <get diagnostics statement> is used to obtain information from an occupied condition area, referenced by its ordinal position within the first diagnostics area.

4.31 Host languages

This part of ISO/IEC 9075 specifies the actions of <externally-invoked procedure>s in SQL-client modules when those <externally-invoked procedure>s are called by programs written in certain specified programming languages, called *host languages*. The term “*PLN* program”, where *PLN* is the name of a host language, refers to a program that conforms to the specification for that host language as identified in Clause 2, “Normative references”.

This part of ISO/IEC 9075 specifies a mechanism whereby SQL language may be embedded in programs that otherwise conform to any of the same specified programming language standards.

NOTE 64 — Interfaces between SQL and the Java programming language are defined in [ISO9075-10] and [ISO9075-13].

Although there are obvious mappings between many SQL data types and the data types of most host languages, this is not the case for all SQL data types or for all host languages.

For the purposes of interfacing with programming languages, the data types DATE, TIME, TIMESTAMP, and INTERVAL shall be converted to or from character strings in those programming languages by means of a <cast specification>. It is anticipated that future evolution of programming language standards will support data types corresponding to these four SQL data types; this standard will then be amended to reflect the availability of those corresponding data types.

The data types CHARACTER, CHARACTER VARYING, and CHARACTER LARGE OBJECT are also mapped to character strings in the programming languages. However, because the facilities available in the programming languages do not provide the same capabilities as those available in SQL, there shall be agreement between the host program and SQL regarding the specific format of the character data being exchanged. Specific syntax for this agreement is provided in this part of ISO/IEC 9075.

For standard programming languages C and COBOL, BOOLEAN values are mapped to integer variables in the host language. For standard programming languages Ada, Fortran, Pascal, and PL/I, BOOLEAN variables are directly supported.

For the purposes of interfacing with programming languages, the data type ARRAY shall be converted to a locator (see Subclause 4.29.5, “Locators”).

For the purposes of interfacing with programming languages, the data type MULTiset shall be converted to a locator (see Subclause 4.29.5, “Locators”).

For the purposes of interfacing with programming languages, user-defined types shall be handled with a locator (see Subclause 4.29.5, “Locators”) or transformed to another SQL data type that has a defined mapping to the host language (see Subclause 4.7.7, “Transforms for user-defined types”).

4.32 Cursors

4.32.1 General description of cursors

This Subclause is modified by Subclause 4.6.1, “General description of cursors”, in ISO/IEC 9075-3.

This Subclause is modified by Subclause 4.9.1, “General description of cursors”, in ISO/IEC 9075-4.

A cursor is a mechanism by which the rows of a table may be acted on (*e.g.*, returned to a host programming language) one at a time.

04 A cursor is specified by a <declare cursor>, a <dynamic declare cursor>, or an <allocate cursor statement>. A cursor specified by a <declare cursor> is a *standing cursor*. A cursor specified by a <dynamic declare cursor> is a *declared dynamic cursor*. A cursor specified by an <allocate cursor statement> is an *extended dynamic cursor*. An extended dynamic cursor whose <allocate cursor statement> contains a <statement cursor> is an *extended statement cursor*; otherwise, the extended dynamic cursor contains a <result set cursor> and is called a *received cursor*. An extended dynamic cursor is said to be *global* if its extended name is global; otherwise, the extended dynamic cursor is *local*. A *dynamic cursor* is either a declared dynamic cursor or an extended dynamic cursor.

03 04 A *declared cursor* is either a standing cursor or a declared dynamic cursor. A declared cursor has a <cursor name>. A <declare cursor> or a <dynamic declare cursor> is immediately contained in the <module contents> of an <SQL-client module definition>. The scope of a <cursor name> is the innermost <SQL-client module definition> *M* that contains it, with the exception of any <SQL schema statement>s contained in *M*.

Every cursor has an associated cursor declaration. For a declared cursor, the cursor declaration is specified by the cursor's <declare cursor> or <dynamic declare cursor>. For an extended dynamic cursor, the cursor declaration is implicitly created by an <allocate cursor statement>.

03 04 A cursor declaration is described by a cursor declaration descriptor. A *cursor declaration descriptor* includes:

- **03** The kind of cursor (standing, declared dynamic, extended statement, or received).
- **04** The provenance of the cursor:
 - If the cursor is a declared cursor or a local extended dynamic cursor, then an indication of an SQL-client module.

- Otherwise, an SQL-session identifier.
- 03 The name of the cursor:
- If the cursor is a declared cursor, then a <cursor name>.
 - 03 Otherwise, an extended name and its scope (GLOBAL or LOCAL).
- 03 The cursor's origin:
- If the cursor is a standing cursor, then the <cursor specification> contained in the <declare cursor>.
 - If the cursor is a declared dynamic cursor, then the <statement name> contained in the <dynamic declare cursor>.
 - If the cursor is an extended statement cursor, then the prepared statement specified by the extended statement name.
 - 03 If the cursor is a received cursor, then the cursor's <specific routine designator>.
- The cursor's declared properties:
- The cursor's declared sensitivity property (either SENSITIVE, INSENSITIVE, or ASENSITIVE).
 - The cursor's declared scrollability property (either SCROLL or NO SCROLL).
 - The cursor's declared holdability property (either WITH HOLD or WITHOUT HOLD).
 - The cursor's declared returnability property (either WITH RETURN or WITHOUT RETURN).

A cursor declaration descriptor is identified by the combination of its kind, its provenance, and its name.

A cursor is described by a cursor instance descriptor. A *cursor instance descriptor* includes:

- A cursor declaration descriptor.
- An SQL-session identifier.
- The cursor's state (either open or closed).
- If the cursor is open, then the cursor's result set descriptor.

A cursor instance descriptor is identified by its cursor declaration descriptor and its SQL-session identifier.

The term “cursor”, unqualified by either “declaration” or “instance”, refers to a cursor instance descriptor.

For every <declare cursor> in an <SQL-client module definition>, a cursor declaration descriptor is included in the SQL-client module, and is effectively created and destroyed at the same time as the SQL-client module. A cursor instance descriptor for the standing cursor is effectively created when an externally-invoked procedure of that SQL-client module is invoked. The cursor instance descriptor is destroyed at the end of the SQL-session.

For every <dynamic declare cursor> in an <SQL-client module definition>, a cursor declaration descriptor is included in the SQL-client module, and is effectively created and destroyed at the same time as the SQL-client module. A cursor instance descriptor for the declared dynamic cursor is effectively created when an externally-invoked procedure of that SQL-client module is invoked, and destroyed when a <deallocate prepared statement> is executed that deallocates the prepared statement on which the declared dynamic cursor is based. The cursor instance descriptor is also destroyed at the end of the SQL-session (if not earlier).

An extended statement cursor is effectively created (*i.e.*, both the cursor declaration descriptor and the cursor instance descriptor) when an <allocate cursor statement> containing a <statement cursor> is executed within an SQL-session and destroyed when that SQL-session is terminated or when a <deallocate prepared statement> is executed that deallocates the prepared statement on which the dynamic cursor is based.

A received cursor is effectively created (*i.e.*, both the cursor declaration descriptor and the cursor instance descriptor) when an <allocate cursor statement> containing a <result set cursor> is executed within an SQL-session and destroyed when that SQL-session is terminated.

A cursor instance descriptor is in either the open state or the closed state. The initial state of a cursor is the closed state. A standing cursor is placed in the open state by an <open statement> and returned to the closed state by a <close statement> or a <rollback statement>. A declared dynamic cursor is placed in the open state by a <dynamic open statement> and returned to the closed state by a <dynamic close statement>. An extended statement cursor is placed in the open state by an <dynamic open statement> and returned to the closed state by a <dynamic close statement>. A received cursor is placed in the open state by an <allocate cursor statement>, and is advanced to the next result set in the result set sequence by a <dynamic close statement>. After the last result set in the result set sequence has been processed, the received cursor is placed in the closed state by a <dynamic close statement>. A received cursor is also placed in the closed state if the associated SQL-invoked procedure is invoked again while the received cursor is open. An open cursor that was not defined as a holdable cursor is also closed by a <commit statement>. (In the case of a received cursor, if the current result set of the received cursor is not holdable, then a <commit statement> causes the received cursor to open on the next result set in the result set sequence, if any.)

A cursor in the open state identifies a result set and a position relative to the ordering of that result set. If the cursor's <cursor specification> does not simply contain an <order by clause>, or simply contains an <order by clause> that does not specify the order of the rows completely, then the rows of the result set have an order that is defined only to the extent that the <order by clause> specifies an order and is otherwise implementation-dependent.

NOTE 65 — A definition of “result set” is given in [Subclause 3.1.6, “Definitions provided in Part 2”](#).

When the ordering of a cursor is not defined by an <order by clause>, the relative position of two rows is implementation-dependent. When the ordering of a cursor is partially determined by an <order by clause>, then the relative positions of two rows are determined only by the <order by clause>; if the two rows have equal values for the purpose of evaluating the <order by clause>, then their relative positions are implementation-dependent.

A cursor whose result set's <cursor specification> simply contains an <order by clause> is an *ordered cursor*.

A cursor is either *updatable* or *not updatable*. A received cursor is not updatable. For other kinds of cursors, if FOR UPDATE OF is contained in the <cursor specification> of the result set descriptor of the cursor, or if the <cursor specification> CS identified by the cursor is a <query specification> that is simply updatable, FOR READ ONLY is not contained in CS, the cursor is not ordered, and the cursor is not scrollable, then the cursor is *updatable*; otherwise, the cursor is *not updatable*. The operations of update and delete are permitted for updatable cursors, subject to constraining Access Rules.

The position of a cursor in the open state is either before a certain row, on a certain row, or after the last row. If a cursor is on a row, then that row is the current row of the cursor. A cursor may be before the first row or after the last row of a result set even though the result set's sequence of rows is empty. When a cursor is initially opened, the position of the cursor is before the first row.

A cursor declaration descriptor and a result set descriptor have four properties: the sensitivity property (either SENSITIVE, INSENSITIVE, or ASENSITIVE), the scrollability property (either SCROLL or NO SCROLL), the holdability property (either WITH HOLD or WITHOUT HOLD), and the returnability property (either WITH RETURN or WITHOUT RETURN). The four declared cursor properties are set to the explicit or implicit

<cursor sensitivity>, <cursor scrollability>, <cursor holdability>, and <cursor returnability>, respectively, of the <cursor properties> contained in the <declare cursor>, <dynamic declare cursor>, or <allocate cursor statement>.

When a cursor is opened, the four declared cursor properties are copied to the result set descriptor's operational properties. In the case of a declared dynamic cursor or an extended statement cursor, the operational properties are modified as specified by the explicit <cursor attributes> of the prepared statement that is the <cursor specification> of the cursor. In the case of a received cursor, the operational sensitivity property, the operational scrollability property, and the operational holdability property are preserved from their values as established by the SQL-invoked procedure that created the result set, and the operational returnability property is implementation-defined.

A cursor whose operational sensitivity property is SENSITIVE is said to be *sensitive*; a cursor whose operational sensitivity property is INSENSITIVE is said to be *insensitive*; and a cursor whose operational sensitivity property is ASENSITIVE is said to be *asensitive*.

A cursor whose operational scrollability property is SCROLL is said to be *scrollable*.

A cursor whose operational holdability property is WITH HOLD is said to be *holdable*.

A cursor whose operational returnability property is WITH RETURN is said to be a *with-return* cursor.

A holdable cursor that has been held open retains its position when the new SQL-transaction is initiated. However, if the holdable cursor is a standing cursor, then before either an <update statement: positioned> or a <delete statement: positioned> is permitted to reference that standing cursor in the new SQL-transaction, a <fetch statement> must be issued against the standing cursor (otherwise, an exception condition is raised). If the holdable cursor is a dynamic cursor, then before a <dynamic update statement: positioned>, a <dynamic delete statement: positioned>, a <preparable dynamic update statement: positioned>, or a <preparable dynamic delete statement: positioned> is permitted to reference that dynamic cursor in the new SQL-transaction, a <dynamic fetch statement> must be issued against the dynamic cursor (otherwise an exception condition is raised).

The following aspects of a standing cursor can be determined from the <cursor specification> in its cursor declaration descriptor: whether it is ordered, whether it is updatable, the underlying tables, its sensitivity property, its scrollability property, its holdability property, and its returnability property. Accordingly, terminology regarding these aspects of a standing cursor may also be applied to the cursor declaration descriptor, for example, in the Syntax Rules, Access Rules, and Conformance Rules.

4.32.2 Operations on and using cursors

A <fetch statement> positions an open standing cursor on a specified row of the standing cursor's ordering and retrieves the values of the columns of that row. An <update statement: positioned> updates the current row of the standing cursor. A <delete statement: positioned> deletes the current row of the standing cursor.

A <dynamic fetch statement> positions an open dynamic cursor on a specified row of the dynamic cursor's ordering and retrieves the values of the columns of that row. A <dynamic update statement: positioned> updates the current row of the dynamic cursor. A <dynamic delete statement: positioned> deletes the current row of the dynamic cursor. A <preparable dynamic delete statement: positioned> is used to delete rows through a dynamic cursor when the precise format of the statement is not known until runtime. A <preparable dynamic update statement: positioned> is used to update rows through a dynamic cursor when the precise format of the statement is not known until runtime.

A *positioned delete statement* is an SQL-statement that is a <delete statement: positioned>, a <dynamic delete statement: positioned>, or a <preparable dynamic delete statement: positioned>.

A *positioned update statement* is an SQL-statement that is a <update statement: positioned>, a <dynamic update statement: positioned>, or a <preparable dynamic update statement: positioned>.

If an error occurs during the execution of an SQL-statement that identifies a cursor, then, except where otherwise explicitly defined, the effect, if any, on the position or state of that cursor is implementation-dependent.

If a completion condition is raised during the execution of an SQL-statement that identifies a cursor, then the particular SQL-statement identifying that open cursor on which the completion condition is returned is implementation-dependent.

The following paragraphs define several terms used to discuss issues relating to a cursor's operational sensitivity property:

A change to SQL-data is said to be *independent* of a cursor *CR* if and only if it is not made by an <update statement: positioned> or a <delete statement: positioned> that is positioned on *CR*.

A change to SQL-data is said to be *significant* to *CR* if and only if it is independent of *CR*, and, had it been committed before *CR* was opened, would have caused the sequence of rows in the result set descriptor of *CR* to be different in any respect.

A change to SQL-data is said to be *visible* to *CR* if and only if it has an effect on the sequence of rows *SR* of the result set descriptor of *CR* by inserting a row in *SR*, deleting a row from *SR*, changing the value of a column of a row of *SR*, or reordering the rows of *SR*.

If a cursor is open, and the SQL-transaction in which the cursor was opened makes a significant change to SQL-data, then whether that change is visible through that cursor before it is closed is determined as follows:

- If the cursor is insensitive, then significant changes are not visible.
- If the cursor is sensitive, then significant changes are visible.
- If the cursor is asensitive, then the visibility of significant changes is implementation-dependent.

If a holdable cursor is open during an SQL-transaction *T* and it is held open for a subsequent SQL-transaction, then whether any significant changes made to SQL-data (by *T* or any subsequent SQL-transaction in which the cursor is held open) are visible through that cursor in the subsequent SQL-transaction before that cursor is closed is determined as follows:

- If the cursor is insensitive, then significant changes are not visible.
- If the cursor is sensitive, then the visibility of significant changes is implementation-defined.
- If the cursor is asensitive, then the visibility of significant changes is implementation-dependent.

If *CR* is a with-return cursor, then the <cursor specification> *CR* included in the result set descriptor of *CR* defines a returned result set. A with-return cursor, if declared in an SQL-invoked procedure and in the open state when the procedure returns to its invoker, yields a returned result set that can be accessed by the invoker of the procedure that generates it.

NOTE 66 — Definitions of “returned result set” and “with-return cursor” are given in [Subclause 3.1.6, “Definitions provided in Part 2”](#).

4.33 SQL-statements

4.33.1 Classes of SQL-statements

This Subclause is modified by Subclause 4.10.1, “Classes of SQL-statements”, in ISO/IEC 9075-4.

An SQL-statement is a string of characters that conforms to the Format and Syntax Rules specified in the parts of ISO/IEC 9075. Most SQL-statements can be prepared for execution and executed in an SQL-client module, in which case they are prepared when the SQL-client module is created and executed when the containing externally-invoked procedure is called (see Subclause 4.22, “SQL-client modules”).

Most SQL-statements can be prepared for execution and executed in additional ways. These are:

- In an embedded SQL host program, in which case they are prepared when the embedded SQL host program is preprocessed (see Subclause 4.23, “Embedded syntax”).
- Being prepared and executed by the use of SQL-dynamic statements (which are themselves executed in an SQL-client module or an embedded SQL host program—see Subclause 4.24, “Dynamic SQL concepts”).
- Direct invocation, in which case they are effectively prepared immediately prior to execution (see Subclause 4.25, “Direct invocation of SQL”).

In this part of ISO/IEC 9075, there are at least six ways of classifying SQL-statements:

- According to their effect on SQL objects, whether persistent objects, *i.e.*, SQL-data, SQL-client modules, and schemas, or transient objects, such as SQL-sessions and other SQL-statements.
- According to whether or not they start an SQL-transaction, or can, or shall, be executed when no SQL-transaction is active.
- According to whether they possibly read SQL-data or possibly modify SQL-data.
- According to whether or not they may be embedded.
- According to whether they may be dynamically prepared and executed.
- According to whether or not they may be directly executed.

This part of ISO/IEC 9075 permits SQL-implementations to provide additional, implementation-defined, statements that may fall into any of these categories. This Subclause will not mention those statements again, as their classification is implementation-defined.

The main classes of SQL-statements are:

- SQL-schema statements; these may have a persistent effect on the set of schemas.
- SQL-data statements; some of these, the SQL-data change statements, may have a persistent effect on SQL-data.
- SQL-transaction statements; except for the <commit statement>, these, and the following classes, have no effects that persist when an SQL-session is terminated.
- SQL-control statements.
- SQL-connection statements.

- SQL-session statements.
- SQL-diagnostics statements.
- SQL-dynamic statements.
- SQL embedded exception declaration.

4.33.2 SQL-statements classified by function

4.33.2.1 SQL-schema statements

*This Subclause is modified by Subclause 4.10.2.1, “SQL-schema statements”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 4.13.1.1, “SQL-schema statements”, in ISO/IEC 9075-9.*

The following are the SQL-schema statements:

- <schema definition>.
- <drop schema statement>.
- <domain definition>.
- <drop domain statement>.
- <table definition>.
- <drop table statement>.
- <view definition>.
- <drop view statement>.
- <assertion definition>.
- <drop assertion statement>.
- <alter table statement>.
- <alter domain statement>.
- <grant privilege statement>.
- <revoke statement>.
- <character set definition>.
- <drop character set statement>.
- <collation definition>.
- <drop collation statement>.
- <transliteration definition>.

- <drop transliteration statement>.
- <trigger definition>.
- <drop trigger statement>.
- <user-defined type definition>.
- <alter type statement>.
- <drop data type statement>.
- <user-defined ordering definition>.
- <drop user-defined ordering statement>.
- <user-defined cast definition>.
- <drop user-defined cast statement>.
- <transform definition>.
- <alter transform statement>.
- <drop transform statement>.
- <schema routine>.
- <alter routine statement>.
- <drop routine statement>.
- <sequence generator definition>.
- <alter sequence generator statement>.
- <drop sequence generator statement>.
- <role definition>.
- <grant role statement>.
- <drop role statement>.

4.33.2.2 SQL-data statements

The following are the SQL-data statements:

- <temporary table declaration>.
- <declare cursor>.
- <open statement>.
- <close statement>.
- <fetch statement>.

- <select statement: single row>.
- <free locator statement>.
- <hold locator statement>.
- <dynamic declare cursor>.
- <allocate cursor statement>.
- <dynamic select statement>.
- <dynamic open statement>.
- <dynamic close statement>.
- <dynamic fetch statement>.
- <direct select statement: multiple rows>.
- <dynamic single row select statement>.
- All SQL-data change statements.

4.33.2.3 SQL-data change statements

The following are the SQL-data change statements:

- <insert statement>.
- <delete statement: searched>.
- <delete statement: positioned>.
- <update statement: searched>.
- <update statement: positioned>.
- <merge statement>.
- <truncate table statement>
- <dynamic delete statement: positioned>.
- <preparable dynamic delete statement: positioned>.
- <dynamic update statement: positioned>.
- <preparable dynamic update statement: positioned>.

4.33.2.4 SQL-transaction statements

The following are the SQL-transaction statements:

- <start transaction statement>.

4.33 SQL-statements

- <set transaction statement>.
- <set constraints mode statement>.
- <commit statement>.
- <rollback statement>.
- <savepoint statement>.
- <release savepoint statement>.

4.33.2.5 SQL-connection statements

The following are the SQL-connection statements:

- <connect statement>.
- <set connection statement>.
- <disconnect statement>.

4.33.2.6 SQL-control statements

This Subclause is modified by Subclause 4.10.2.2, “SQL-control statements”, in ISO/IEC 9075-4.

The following are the SQL-control statements:

- <call statement>.
- <return statement>.

4.33.2.7 SQL-session statements

This Subclause is modified by Subclause 4.13.1.2, “SQL-session statements”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 4.6.1.1, “SQL-session statements”, in ISO/IEC 9075-14.

The following are the SQL-session statements:

- <set session characteristics statement>.
- <set session user identifier statement>.
- <set role statement>.
- <set local time zone statement>.
- <set catalog statement>.
- <set schema statement>.

- <set names statement>.
- <set path statement>.
- <set transform group statement>.
- <set session collation statement>.

4.33.2.8 SQL-diagnostics statements

This Subclause is modified by Subclause 4.10.2.4, “SQL-diagnostics statements”, in ISO/IEC 9075-4.

The following are the SQL-diagnostics statements:

- <get diagnostics statement>.

4.33.2.9 SQL-dynamic statements

The following are the SQL-dynamic statements:

- <execute immediate statement>.
- <allocate descriptor statement>.
- <deallocate descriptor statement>.
- <get descriptor statement>.
- <set descriptor statement>.
- <prepare statement>.
- <deallocate prepared statement>.
- <describe input statement>.
- <describe output statement>.
- <execute statement>.

4.33.2.10 SQL embedded exception declaration

The following is the SQL embedded exception declaration:

- <embedded exception declaration>.

4.33.3 SQL-statements and SQL-data access indication

Some SQL-statements may be classified either as SQL-statements that *possibly read SQL-data* or that *possibly modify SQL-data*. A given SQL-statement belongs to at most one such class.

The following SQL-statements possibly read SQL-data:

- SQL-data statements other than SQL-data change statements, <free locator statement>, and <hold locator statement>.
- SQL-statements that contain a <query expression> and are not SQL-statements that possibly modify SQL-data.

The following SQL-statements possibly modify SQL-data:

- SQL-schema statements.
- SQL-data change statements.

4.33.4 SQL-statements and transaction states

This Subclause is modified by Subclause 4.10.3, “SQL-statements and transaction states”, in ISO/IEC 9075-4.

The following SQL-statements are transaction-initiating SQL-statements, *i.e.*, if there is no current SQL-transaction, and an SQL-statement of this class is executed, then an SQL-transaction is initiated, usually before execution of that SQL-statement proceeds:

- All SQL-schema statements
- The following SQL-transaction statements:
 - <start transaction statement>.
 - <savepoint statement>.
 - <commit statement>.
 - <rollback statement>.
- The following SQL-data statements:
 - <open statement>.
 - <close statement>.
 - <fetch statement>.
 - <select statement: single row>.
 - <insert statement>.
 - <delete statement: searched>.
 - <delete statement: positioned>.
 - <update statement: searched>.

- <update statement: positioned>.
- <merge statement>.
- <truncate table statement>
- <allocate cursor statement>.
- <dynamic open statement>.
- <dynamic close statement>.
- <dynamic fetch statement>.
- <direct select statement: multiple rows>.
- <dynamic single row select statement>.
- <dynamic delete statement: positioned>.
- <preparable dynamic delete statement: positioned>.
- <dynamic update statement: positioned>.
- <preparable dynamic update statement: positioned>.
- <free locator statement>.
- <hold locator statement>.

— The following SQL-dynamic statements:

- <describe input statement>.
- <describe output statement>.
- <allocate descriptor statement>.
- <deallocate descriptor statement>.
- <get descriptor statement>.
- <set descriptor statement>.
- <prepare statement>.
- <deallocate prepared statement>.

With the exception of <start transaction statement>, every transaction-initiating SQL-statement implicitly initiates an SQL-transaction if there is no current SQL-transaction when it is executed, in which case the execution of the initiating SQL-statement is included in the initiated SQL-transaction. In the case of <start transaction statement>, SQL-transaction initiation is the primary effect of executing the SQL-statement itself.

Whether or not an <execute immediate statement> starts a transaction depends on the content of the <SQL statement variable> referenced by the <execute immediate statement> at the time it is executed. Whether or not an <execute statement> starts a transaction depends on the content of the <SQL statement variable> referenced by the <prepare statement> at the time the prepared statement referenced by the <execute statement> was prepared. In both cases, if the content of the <SQL statement variable> was a transaction-initiating SQL-statement, then the <execute immediate statement> or <execute statement> is treated as a transaction-initiating statement; otherwise, it is not treated as a transaction-initiating statement.

The following SQL-statements are not transaction-initiating SQL-statements, *i.e.*, then executing an SQL-statement of this class does not affect the existence or absence of an SQL-transaction.

- All SQL-transaction statements except <start transaction statement>s, <savepoint statement>s, <commit statement>s, and <rollback statement>s.
- All SQL-connection statements.
- All SQL-session statements.
- All SQL-diagnostics statements.
- SQL embedded exception declarations.
- The following SQL-data statements:
 - <temporary table declaration>.
 - <declare cursor>.
 - <dynamic declare cursor>.
 - <dynamic select statement>.

The following SQL-statements are possibly transaction-initiating SQL-statements:

- <return statement>.
- <call statement>.

If an <SQL control statement> causes the evaluation of a <query expression> and there is no current SQL-transaction, then an SQL-transaction is initiated before evaluation of the <query expression>.

4.33.5 SQL-statement atomicity and statement execution contexts

This Subclause is modified by Subclause 4.10.4, “SQL-statement atomicity and statement execution contexts”, in ISO/IEC 9075-4.

The execution of all SQL-statements other than certain SQL-control statements and certain SQL-transaction statements is atomic with respect to recovery. Such an SQL-statement is called an *atomic SQL-statement*. An SQL-statement that is not an atomic SQL-statement is called a *non-atomic SQL statement*.

The following are non-atomic SQL-statements:

- <call statement>
- <execute statement>
- <execute immediate statement>
- <commit statement>
- <return statement>
- <rollback statement>
- <savepoint statement>

All other SQL-statements are atomic SQL-statements.

A statement execution context is either *atomic* or *non-atomic*.

The statement execution context brought into existence by the execution of a non-atomic SQL-statement is a *non-atomic execution context*.

The statement execution context brought into existence by the execution of an atomic SQL-statement is an *atomic execution context*.

Within one execution context, another execution context may become active. This latter execution context is said to be a *more recent execution context* than the former. If there is no execution context that is more recent than execution context *EC*, then *EC* is said to be the *most recent execution context*.

If there is no atomic execution context that is more recent than atomic execution context *AEC*, then *AEC* is the *most recent atomic execution context*.

An SQL-transaction cannot be explicitly terminated within an atomic execution context. If the execution of an atomic SQL-statement is unsuccessful, then the changes to SQL-data or schemas made by the SQL-statement are canceled.

4.33.6 Embeddable SQL-statements

This Subclause is modified by Subclause 4.10.5, “Embeddable SQL-statements”, in ISO/IEC 9075-4.

The following SQL-statements are embeddable in an embedded SQL host program, and may be the <SQL procedure statement> in an <externally-invoked procedure> in an <SQL-client module definition>:

- All SQL-schema statements.
- All SQL-transaction statements.
- All SQL-connection statements.
- All SQL-session statements.
- All SQL-dynamic statements.
- All SQL-diagnostics statements.
- The following SQL-data statements:
 - <allocate cursor statement>.
 - <open statement>.
 - <dynamic open statement>.
 - <close statement>.
 - <dynamic close statement>.
 - <fetch statement>.
 - <dynamic fetch statement>.
 - <select statement: single row>.

- <insert statement>.
- <delete statement: searched>.
- <delete statement: positioned>.
- <dynamic delete statement: positioned>.
- <update statement: searched>.
- <update statement: positioned>.
- <merge statement>.
- <truncate table statement>
- <dynamic update statement: positioned>.
- <hold locator statement>.
- <free locator statement>.

— The following SQL-control statements:

- <call statement>.
- <return statement>.

The following SQL-statements are embeddable in an embedded SQL host program, and may occur in an <SQL-client module definition>, though not in an <externally-invoked procedure>:

- <temporary table declaration>.
- <declare cursor>.
- <dynamic declare cursor>.

The following SQL-statements are embeddable in an embedded SQL host program, but may not occur in an <SQL-client module definition>:

- SQL embedded exception declarations.

Consequently, the following SQL-data statements are not embeddable in an embedded SQL host program, nor may they occur in an <SQL-client module definition>, nor be the <SQL procedure statement> in an <externally-invoked procedure> in an <SQL-client module definition>:

- <dynamic select statement>.
- <dynamic single row select statement>.
- <direct select statement: multiple rows>.
- <preparable dynamic delete statement: positioned>.
- <preparable dynamic update statement: positioned>.

4.33.7 Preparable and immediately executable SQL-statements

This Subclause is modified by Subclause 4.10.6, “Preparable and immediately executable SQL-statements”, in ISO/IEC 9075-4.

The following SQL-statements are preparable:

- All SQL-schema statements.
- All SQL-transaction statements.
- All SQL-session statements.
- The following SQL-data statements:
 - <delete statement: searched>.
 - <dynamic select statement>.
 - <dynamic single row select statement>.
 - <insert statement>.
 - <update statement: searched>.
 - <truncate table statement>
 - <merge statement>.
 - <preparable dynamic delete statement: positioned>.
 - <preparable dynamic update statement: positioned>.
 - <preparable implementation-defined statement>.
 - <hold locator statement>.
 - <free locator statement>.
- The following SQL-control statements:
 - <call statement>.

Consequently, the following SQL-statements are not preparable:

- All SQL-connection statements.
- All SQL-dynamic statements.
- All SQL-diagnostics statements.
- SQL embedded exception declarations.
- The following SQL-data statements:
 - <allocate cursor statement>.
 - <open statement>.
 - <dynamic open statement>.

4.33 SQL-statements

- <close statement>.
- <dynamic close statement>.
- <fetch statement>.
- <dynamic fetch statement>.
- <select statement: single row>.
- <delete statement: positioned>.
- <dynamic delete statement: positioned>.
- <update statement: positioned>.
- <dynamic update statement: positioned>.
- <direct select statement: multiple rows>.
- <temporary table declaration>.
- <declare cursor>.
- <dynamic declare cursor>.

— The following SQL-control statements:

- <return statement>.

Any preparable SQL-statement can be executed immediately, with the exception of:

- <dynamic select statement>.
- <dynamic single row select statement>.

4.33.8 Directly executable SQL-statements

This Subclause is modified by Subclause 4.10.7, “Directly executable SQL-statements”, in ISO/IEC 9075-4.

The following SQL-statements may be executed directly:

- All SQL-schema statements.
- All SQL-transaction statements.
- All SQL-connection statements.
- All SQL-session statements.
- The following SQL-data statements:
 - <temporary table declaration>.
 - <direct select statement: multiple rows>.
 - <insert statement>.

- <delete statement: searched>.
- <update statement: searched>.
- <truncate table statement>
- <merge statement>.

— The following SQL-control statements:

- <call statement>.
- <return statement>.

Consequently, the following SQL-statements may not be executed directly:

— All SQL-dynamic statements.

— All SQL-diagnostics statements.

— SQL embedded exception declarations.

— The following SQL-data statements:

- <declare cursor>.
- <dynamic declare cursor>.
- <allocate cursor statement>.
- <open statement>.
- <dynamic open statement>.
- <close statement>.
- <dynamic close statement>.
- <fetch statement>.
- <dynamic fetch statement>.
- <select statement: single row>.
- <dynamic select statement>.
- <dynamic single row select statement>.
- <delete statement: positioned>.
- <dynamic delete statement: positioned>.
- <preparable dynamic delete statement: positioned>.
- <update statement: positioned>.
- <dynamic update statement: positioned>.
- <preparable dynamic update statement: positioned>.

— <free locator statement>.

— <hold locator statement>.

4.34 Basic security model

4.34.1 Authorization identifiers

An authorization identifier identifies a set of privileges. An authorization identifier is either a user identifier or a role name. A user identifier represents a user of the SQL-implementation. Any mapping of user identifiers to operating system users is implementation-defined. A role name represents a role.

4.34.1.1 SQL-session authorization identifiers

An SQL-session has a user identifier called the *SQL-session user identifier*. When an SQL-session is initiated, the SQL-session user identifier is determined in an implementation-defined manner, unless the session is initiated using a <connect statement>. The value of the SQL-session user identifier can never be the null value. The SQL-session user identifier can be determined by using SESSION_USER.

An SQL-session context contains a time-varying sequence of cells, known as the *authorization stack*, each cell of which contains either a user identifier, a role name, or both. This stack is maintained using a “last-in, first-out” discipline, and effectively only the top cell is visible. When an SQL-session is started, by explicit or implicit execution of a <connect statement>, the authorization stack is initialized with one cell, which contains only the user identifier known as the *SQL-session user identifier*; a role name, known as the *SQL-session role name* may be added subsequently.

Let *E* be an externally-invoked procedure, SQL-invoked routine, triggered action, prepared statement, or directly executed statement. When *E* is invoked, a copy of the top cell is pushed onto the authorization stack. If the invocation of *E* is to be under definer's rights, then the contents of the top cell are replaced with the authorization identifier of the owner of *E*. On completion of the execution of *E*, the top cell is removed.

The contents of the top cell in the authorization stack of the current SQL-session context determine the privileges for the execution of each SQL-statement. The user identifier, if any, in this cell is known as the *current user identifier*; the role name, if any, is known as the *current role name*. They may be determined using CURRENT_USER and CURRENT_ROLE, respectively.

At a given time, there may be no current user identifier or no current role name, but at least one or the other is always present.

NOTE 67 — The privileges granted to PUBLIC are available to all of the <authorization identifier>s in the SQL-environment.

The <set session user identifier statement> changes the value of the current user identifier and of the SQL-session user identifier. The <set role statement> changes the value of the current role name.

The term *current authorization identifier* denotes an authorization identifier in the top cell of the authorization stack.

4.34.1.2 SQL-client module authorization identifiers

If an <SQL-client module definition> contains a <module authorization identifier> *MAI*, then *MAI* is the owner of the corresponding SQL-client module *M* and is used as the current authorization identifier for the execution of each externally-invoked procedure in *M*. If *M* has no owner, then the current user identifier and the current role name of the SQL-session are used as the current user identifier and current role name, respectively, for the execution of each externally-invoked procedure in *M*.

4.34.1.3 SQL-schema authorization identifiers

Every schema has an owner, determined at the time of its creation from a <schema definition> *SD*. That owner is

Case:

- If *SD* simply contains a <schema authorization identifier> *SAI*, then *SAI*.
- If *SD* is simply contained in an <SQL-client module definition> that contains a <module authorization identifier> *MAI*, then *MAI*.
- Otherwise, the SQL-session user identifier.

4.34.2 Privileges

This Subclause is modified by Subclause 4.11.1, “Privileges”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 4.14.1, “Privileges”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 4.10, “Privileges”, in ISO/IEC 9075-13.

This Subclause is modified by Subclause 4.7.1, “Privileges”, in ISO/IEC 9075-14.

04091314 A privilege authorizes a given category of <action> to be performed by a specified <authorization identifier> on a specified object. The object in question is one of the following:

- A base table.
- A view.
- A view component.
- A column.
- A domain.
- A character set.
- A collation.
- A transliteration.
- A user-defined type.
- A table/method pair.

- An SQL-invoked routine.
- A sequence generator.

Each privilege is represented by a *privilege descriptor*. A privilege descriptor contains:

- The identification of the object on which the privilege is granted.
- The <authorization identifier> of the grantor of the privilege.
- The <authorization identifier> of the grantee of the privilege.
- Identification of the <action> that the privilege allows.
- An indication of whether or not the privilege is grantable.
- An indication of whether or not the privilege has the WITH HIERARCHY OPTION specified.

The <action>s that can be specified are:

- INSERT.
- INSERT (<column name list>).
- UPDATE.
- UPDATE (<column name list>).
- DELETE.
- SELECT.
- SELECT (<column name list>).
- SELECT (<privilege method list>).
- REFERENCES.
- REFERENCES (<column name list>).
- USAGE.
- UNDER.
- TRIGGER.
- EXECUTE.

A privilege descriptor with an <action> of INSERT, UPDATE, DELETE, SELECT, TRIGGER, or REFERENCES is called a *table privilege descriptor* and identifies the existence of a privilege on the table or view component identified by the privilege descriptor. If a table privilege descriptor identifies a view component, the privilege descriptor is called a *view component table privilege descriptor*.

A privilege descriptor with an <action> of SELECT (<column name list>), INSERT (<column name list>), UPDATE (<column name list>), or REFERENCES (<column name list>) is called a *column privilege descriptor* and identifies the existence of a privilege on the columns in the table or view component identified by the privilege descriptor. If a column privilege descriptor identifies a view component, the privilege descriptor is called a *view component column privilege descriptor*.

A privilege descriptor with an <action> of SELECT (<privilege method list>) is called a *table/method privilege descriptor* and identifies the existence of a privilege on the methods of the structured type of the table identified by the privilege descriptor.

A table privilege descriptor specifies that the privilege identified by the <action> (unless the <action> is DELETE) is to be automatically granted by the grantor to the grantee on all columns subsequently added to the table.

09 13 14 A privilege descriptor with an <action> of USAGE is called a *usage privilege descriptor* and describes a privilege on an object of one of the following kinds:

- A domain.
- A user-defined type.
- A character set.
- A collation.
- A transliteration.
- A sequence generator.

A privilege descriptor with an <action> of UNDER is called an *under privilege descriptor* and identifies the existence of the privilege on the structured type identified by the privilege descriptor.

A privilege descriptor with an <action> of EXECUTE is called an *execute privilege descriptor* and identifies the existence of a privilege on the SQL-invoked routine identified by the privilege descriptor.

A grantable privilege is a privilege that may be granted by a <grant privilege statement>. The WITH GRANT OPTION clause of a <grant privilege statement> specifies whether the <authorization identifier> identifying the recipient of a privilege (acting as a grantor) may grant it to others.

Privilege descriptors that represent privileges for the owner of an object have a special grantor value, “_SYSTEM”. This value is reflected in the Information Schema for all privileges that apply to the owner of the object.

NOTE 68 — The Information Schema is defined in [ISO9075-11].

A schema that is owned by a given schema <user identifier> or schema <role name> may contain privilege descriptors that describe privileges granted to other <authorization identifier>s (grantees). The granted privileges apply to objects defined in the current schema.

Direct SQL statements are always executed under invoker's rights.

A *view privilege dependency descriptor* is a descriptor that includes two privilege descriptors, called the *supporting privilege descriptor* and the *dependent privilege descriptor*. A view privilege dependency descriptor is a record that an INSERT, UPDATE, or DELETE privilege of a view, or a column of a view, is directly dependent on another privilege.

4.34.3 Roles

A role, identified by a role name, is, like a user, a potential grantee and grantor of privileges and of other roles. Also like a user, a role can additionally own schemas and other objects.

A role is created by executing a <role definition> and destroyed by executing a <drop role statement>.

A role is granted to one or more authorization identifiers by executing a <grant role statement>, thus conferring on the grantees all the privileges of that role. The granting of a role to an authorization identifier *A* is called a *role authorization* (for *A*).

The privileges of a role with role name *R* are the union of the privileges whose grantee is *R* and the sets of privileges for the role names defined by the role authorizations for *R*. Cycles of role authorizations are prohibited.

The WITH ADMIN OPTION clause of the <grant role statement> for role *R* specifies that each grantee may grant *R* to others, revoke *R* from others, and destroy *R*.

Each role authorization is described by a *role authorization descriptor*. A role authorization descriptor includes:

- The role name of the role.
- The authorization identifier of the grantor.
- The authorization identifier of the grantee.
- An indication of whether or not the role authorization is grantable.

4.34.4 Security model definitions

A role *R* is *applicable for* an authorization identifier *A* if there exists a role authorization descriptor whose role name is *R* and whose grantee is PUBLIC, or *A*, or an applicable role for *A*.

A privilege *P* is *applicable for* an authorization identifier *A* if its grantee is PUBLIC, or *A*, or an applicable role for *A*.

NOTE 69 — “applicable for” is a persistent relationship between persistent objects. Thus, it in no way depends on any SQL-session.

An authorization identifier is *enabled* if it is the current user identifier, the current role name, or a role name that is applicable for the current role name.

A privilege *P* is *current* if *P* is applicable for an enabled authorization identifier.

NOTE 70 — “enabled” and “current” apply to (transient) elements of the current SQL-session context.

4.35 SQL-transactions

4.35.1 General description of SQL-transactions

An *SQL-transaction* (transaction) is a sequence of executions of SQL-statements that is atomic with respect to recovery. These operations are performed by one or more compilation units and SQL-client modules. The operations comprising an SQL-transaction may also be performed by the direct invocation of SQL.

It is implementation-defined whether or not the execution of an SQL-data statement is permitted to occur within the same SQL-transaction as the execution of an SQL-schema statement. If it does occur, then the effect on any open cursor or deferred constraint is implementation-defined. There may be additional implementation-defined restrictions, requirements, and conditions. If any such restrictions, requirements, or conditions are

violated, then an implementation-defined exception condition or a completion condition *warning* with an implementation-defined subclass code is raised.

It is implementation-defined whether or not the dynamic execution of an <SQL dynamic data statement> is permitted to occur within the same SQL-transaction as the dynamic execution of an SQL-schema statement. If it does occur, then the effect on any open cursor, prepared dynamic statement, or deferred constraint is implementation-defined. There may be additional implementation-defined restrictions, requirements, and conditions. If any such restrictions, requirements, or conditions are violated, then an implementation-defined exception condition or a completion condition *warning* with an implementation-defined subclass code is raised.

Each SQL-client module that executes an SQL-statement of an SQL-transaction is associated with that SQL-transaction. Each direct invocation of SQL that executes an SQL-statement of an SQL-transaction is associated with that SQL-transaction. An SQL-transaction is initiated when no SQL-transaction is currently active by direct invocation of SQL that results in the execution of a transaction-initiating <direct SQL statement>. An SQL-transaction is initiated when no SQL-transaction is currently active and an <externally-invoked procedure> is called that results in the execution of a *transaction-initiating* SQL-statement. An SQL-transaction is terminated by a <commit statement> or a <rollback statement>. If an SQL-transaction is terminated by successful execution of a <commit statement>, then all changes made to SQL-data or schemas by that SQL-transaction are made persistent and accessible to all concurrent and subsequent SQL-transactions. If an SQL-transaction is terminated by a <rollback statement> or unsuccessful execution of a <commit statement>, then all changes made to SQL-data or schemas by that SQL-transaction are canceled. Committed changes cannot be canceled. If execution of a <commit statement> is attempted, but certain exception conditions are raised, it is unknown whether or not the changes made to SQL-data or schemas by that SQL-transaction are canceled or made persistent.

4.35.2 Savepoints

An SQL-transaction may be partially rolled back by using a savepoint. The savepoint and its <savepoint name> are established within an SQL-transaction when a <savepoint statement> is executed.

An SQL-transaction has one or more *savepoint levels*, exactly one of which is the *current savepoint level*. The savepoint levels of an SQL-transaction are nested, such that when a *new savepoint level NSL is established*, the current savepoint level *CSL* ceases to be current and *NSL* becomes current. When *NSL is destroyed*, *CSL* becomes current again.

A savepoint level exists in an SQL-session *SS* even when no SQL-transaction is active, this savepoint level remaining the current one when an SQL-transaction is initiated in *SS*.

A savepoint *SP* exists at exactly one savepoint level, namely, the savepoint level that is current when *SP* is established.

If a <rollback statement> references a savepoint *SS*, then all changes made to SQL-data or schema subsequent to the establishment of *SS* are canceled, and all savepoints established since *SS* was established are destroyed.

NOTE 71 — The SQL-transaction is effectively restored to its state as it was immediately following the execution of the <savepoint statement> that established *SS*.

Savepoints existing at savepoint level *SPL* are destroyed when *SPL* is destroyed. Savepoint *SS* in the current savepoint level and all savepoints established since *SS* was established are destroyed when a <release savepoint statement> specifying the savepoint name of *SS* is executed. A savepoint may be replaced by another with the same name within a savepoint level by executing a <savepoint statement> that specifies that name.

It is implementation-defined whether or not, or how, a <rollback statement> that references a <savepoint specifier> affects diagnostics area contents, the contents of SQL descriptor areas, and the status of prepared statements.

4.35.3 Properties of SQL-transactions

This Subclause is modified by Subclause 4.15.1, “Properties of SQL-transactions”, in ISO/IEC 9075-9.

An SQL-transaction has a *constraint mode* for each integrity constraint. The constraint mode for an integrity constraint in an SQL-transaction is described in Subclause 4.17, “Integrity constraints”.

09 An SQL-transaction has an *access mode* that is either *read-only* or *read-write*. The access mode may be explicitly set by a <set transaction statement> before the start of an SQL-transaction or by the use of a <start transaction statement> to start an SQL-transaction; otherwise, it is implicitly set to the default access mode for the SQL-session before each SQL-transaction begins. If no <set session characteristics statement> has set the default access mode for the SQL-session, then the default access mode for the SQL-session is *read-write*. The term *read-only* applies only to viewed tables and persistent base tables.

An SQL-transaction has a *condition area limit*, which is a positive integer that specifies the maximum number of conditions that can be placed in any diagnostics area during execution of an SQL-statement in this SQL-transaction.

SQL-transactions initiated by different SQL-agents that access the same SQL-data or schemas and overlap in time are *concurrent SQL-transactions*.

4.35.4 Isolation levels of SQL-transactions

An SQL-transaction has an *isolation level* that is READ UNCOMMITTED, READ COMMITTED, REPEATABLE READ, or SERIALIZABLE. The isolation level of an SQL-transaction defines the degree to which the operations on SQL-data or schemas in that SQL-transaction are affected by the effects of and can affect operations on SQL-data or schemas in concurrent SQL-transactions. The isolation level of an SQL-transaction when any cursor is held open from the previous SQL-transaction within an SQL-session is the isolation level of the previous SQL-transaction by default. If no cursor is held open, or this is the first SQL-transaction within an SQL-session, then the isolation level is SERIALIZABLE by default. The level can be explicitly set by the <set transaction statement> before the start of an SQL-transaction or by the use of a <start transaction statement> to start an SQL-transaction. If it is not explicitly set, then the isolation level is implicitly set to the default isolation level for the SQL-session before each SQL-transaction begins. If no <set session characteristics statement> has set the default isolation level for the SQL-session, then the default isolation level for the SQL-session is SERIALIZABLE.

Execution of a <set transaction statement> is prohibited after the start of an SQL-transaction and before its termination. Execution of a <set transaction statement> before the start of an SQL-transaction sets the access mode, isolation level, and condition area limit for the single SQL-transaction that is started after the execution of that <set transaction statement>. If multiple <set transaction statement>s are executed before the start of an SQL-transaction, the last such statement is the one whose settings are effective for that SQL-transaction; their actions are not cumulative.

The execution of concurrent SQL-transactions at isolation level SERIALIZABLE is guaranteed to be serializable. A serializable execution is defined to be an execution of the operations of concurrently executing SQL-transac-

tions that produces the same effect as some serial execution of those same SQL-transactions. A serial execution is one in which each SQL-transaction executes to completion before the next SQL-transaction begins.

The isolation level specifies the kind of phenomena that can occur during the execution of concurrent SQL-transactions. The following phenomena are possible:

- 1) *P1* (“Dirty read”): SQL-transaction *T1* modifies a row. SQL-transaction *T2* then reads that row before *T1* performs a COMMIT. If *T1* then performs a ROLLBACK, *T2* will have read a row that was never committed and that may thus be considered to have never existed.
- 2) *P2* (“Non-repeatable read”): SQL-transaction *T1* reads a row. SQL-transaction *T2* then modifies or deletes that row and performs a COMMIT. If *T1* then attempts to reread the row, it may receive the modified value or discover that the row has been deleted.
- 3) *P3* (“Phantom”): SQL-transaction *T1* reads the set of rows *N* that satisfy some <search condition>. SQL-transaction *T2* then executes SQL-statements that generate one or more rows that satisfy the <search condition> used by SQL-transaction *T1*. If SQL-transaction *T1* then repeats the initial read with the same <search condition>, it obtains a different collection of rows.

The four isolation levels guarantee that each SQL-transaction will be executed completely or not at all, and that no updates will be lost. The isolation levels are different with respect to phenomena *P1*, *P2*, and *P3*. [Table 8](#), “SQL-transaction isolation levels and the three phenomena” specifies the phenomena that are possible and not possible for a given isolation level.

03 Table 8 — SQL-transaction isolation levels and the three phenomena

Level	<i>P1</i>	<i>P2</i>	<i>P3</i>
READ UNCOMMITTED	Possible	Possible	Possible
READ COMMITTED	Not Possible	Possible	Possible
REPEATABLE READ	Not Possible	Not Possible	Possible
SERIALIZABLE	Not Possible	Not Possible	Not Possible

NOTE 72 — The exclusion of these phenomena for SQL-transactions executing at isolation level SERIALIZABLE is a consequence of the requirement that such transactions be serializable.

Changes made to SQL-data or schemas by an SQL-transaction in an SQL-session may be perceived by that SQL-transaction in that same SQL-session, and by other SQL-transactions, or by that same SQL-transaction in other SQL-sessions, at isolation level READ UNCOMMITTED, but cannot be perceived by other SQL-transactions at isolation level READ COMMITTED, REPEATABLE READ, or SERIALIZABLE until the former SQL-transaction terminates with a <commit statement>.

Regardless of the isolation level of the SQL-transaction, phenomena *P1*, *P2*, and *P3* shall not occur during the implied reading of schema definitions performed on behalf of executing an SQL-statement, the checking of integrity constraints, and the execution of referential actions associated with referential constraints. The schema definitions that are implicitly read are implementation-dependent. This does not affect the explicit reading of rows from tables in the Information Schema, which is done at the isolation level of the SQL-transaction.

NOTE 73 — The Information Schema is defined in [ISO9075-11].

4.35.5 Implicit rollbacks

The execution of a <rollback statement> may be initiated implicitly by an SQL-implementation when it detects the inability to guarantee the serializability of two or more concurrent SQL-transactions. When this error occurs, an exception condition is raised: *transaction rollback — serialization failure*.

The execution of a <rollback statement> may be initiated implicitly by an SQL-implementation when it detects unrecoverable errors. When such an error occurs, an exception condition is raised: *transaction rollback* with an implementation-defined subclass code.

The execution of a <rollback statement> may be initiated implicitly when an SQL-implementation detects the loss of the current SQL-Connection. See [Subclause 4.36, “SQL-connections”](#).

4.35.6 Effects of SQL-statements in an SQL-transaction

The execution of an SQL-statement within an SQL-transaction has no effect on SQL-data or schemas other than the effect stated in the General Rules for that SQL-statement, in the General Rules for [Subclause 11.8, “<referential constraint definition>”](#), in the General Rules for [Subclause 11.40, “<trigger definition>”](#), and in the General Rules for [Subclause 11.51, “<SQL-invoked routine>”](#). Together with serializable execution, this implies that all read operations are repeatable within an SQL-transaction at isolation level `SERIALIZABLE`, except for:

- 1) The effects of changes to SQL-data or schemas and its contents made explicitly by the SQL-transaction itself.
- 2) The effects of differences in SQL parameter values supplied to externally-invoked procedures.
- 3) The effects of references to time-varying system variables such as `CURRENT_DATE` and `CURRENT_USER`.

4.35.7 Encompassing transactions

In some environments (*e.g.*, remote database access), an SQL-transaction can be part of an encompassing transaction that is controlled by an agent other than the SQL-agent. The encompassing transaction may involve different resource managers, the SQL-implementation being just one instance of such a manager. In such environments, an encompassing transaction shall be terminated via that other agent, which in turn interacts with the SQL-implementation via an interface that may be different from SQL (`COMMIT` or `ROLLBACK`), in order to coordinate the orderly termination of the encompassing transaction. When an SQL-transaction is part of an encompassing transaction that is controlled by an agent other than an SQL-agent and a <rollback statement> is initiated implicitly by an SQL-implementation, then the SQL-implementation will interact with that other agent to terminate that encompassing transaction. The specification of the interface between such agents and the SQL-implementation is beyond the scope of this part of ISO/IEC 9075. However, it is important to note that the semantics of an SQL-transaction remain as defined in the following sense:

- When an agent that is different from the SQL-agent requests the SQL-implementation to rollback an SQL-transaction, the General Rules of [Subclause 17.8, “<rollback statement>”](#), are performed.
- When such an agent requests the SQL-implementation to commit an SQL-transaction, the General Rules of [Subclause 17.7, “<commit statement>”](#), are performed. To guarantee orderly termination of the encom-

passing transaction, this commit operation may be processed in several phases not visible to the application; not all the General Rules of [Subclause 17.7](#), “<commit statement>”, need to be executed in a single phase.

However, even in such environments, the SQL-agent interacts directly with the SQL-server to set characteristics (such as *read-only* or *read-write*, isolation level, and constraints mode) that are specific to the SQL-transaction model.

It is implementation-defined whether SQL-transactions that affect more than one SQL-server are supported. If such SQL-transactions are supported, then the part of each SQL-transaction that affects a single SQL-server is called a *branch transaction* or a branch of the SQL-transaction. If such SQL-transactions are supported, then they generally have all the same characteristics (access mode, condition area limit, and isolation level, as well as constraint mode). However, it is possible to alter some characteristics of such an SQL-transaction at one SQL-server by the use of the SET LOCAL TRANSACTION statement; if a SET LOCAL TRANSACTION statement is executed at an SQL-server before any transaction-initiating SQL-statement, then it may set the characteristics of that *branch* of the SQL-transaction at that SQL-server.

The characteristics of a branch of an SQL-transaction are limited by the characteristics of the SQL-transaction as a whole:

- If the SQL-transaction is read-write, then the branch of the SQL-transaction may be read-write or read-only; if the SQL-transaction is read-only, then the branch of the SQL-transaction shall be read-only.
- If the SQL-transaction has an isolation level of READ UNCOMMITTED, then the branch of the SQL-transaction may have an isolation level of READ UNCOMMITTED, READ COMMITTED, REPEATABLE READ, or SERIALIZABLE.

If the SQL-transaction has an isolation level of READ COMMITTED, then the branch of the SQL-transaction shall have an isolation level of READ COMMITTED, REPEATABLE READ, or SERIALIZABLE.

If the SQL-transaction has an isolation level of REPEATABLE READ, then the branch of the SQL-transaction shall have an isolation level of REPEATABLE READ or SERIALIZABLE.

If the SQL-transaction has an isolation level of SERIALIZABLE, then the branch of the SQL-transaction shall have an isolation level of SERIALIZABLE.

- The condition area limit of a branch of an SQL-transaction is always the same as the condition area limit of the SQL-transaction; SET LOCAL TRANSACTION shall not specify a condition area limit.

SQL-transactions that are not part of an encompassing transaction are terminated by the execution of <commit statement>s and <rollback statement>s. If those statements specify AND CHAIN, then they also initiate a new SQL-transaction with the same characteristics as the SQL-transaction that was just terminated, except that the constraint mode of each integrity constraint reverts to its default mode (*deferred* or *immediate*).

4.36 SQL-connections

An *SQL-connection* is an association between an SQL-client and an SQL-server. An SQL-connection may be established and named by a <connect statement>, which identifies the desired SQL-server by means of an <SQL-server name>. A <connection name> is specified as a <simple value specification> whose value is an <identifier>. Two <connection name>s identify the same SQL-connection if their values, with leading and trailing <space>s removed, are equivalent according to the rules for <identifier> comparison in [Subclause 5.2](#), “<token> and <separator>”. It is implementation-defined how an SQL-implementation uses <SQL-server

name> to determine the location, identity, and communication protocol required to access the SQL-server and create an SQL-session.

An SQL-connection is an *active SQL-connection* if any SQL-statement that initiates or requires an SQL-transaction has been executed at its SQL-server via that SQL-connection during the current SQL-transaction.

An SQL-connection is either *current* or *dormant*. If the SQL-connection established by the most recently executed implicit or explicit <connect statement> or <set connection statement> has not been terminated, then that SQL-connection is the *current SQL-connection*; otherwise, there is no current SQL-connection. An existing SQL-connection that is not the current SQL-connection is a *dormant SQL-connection*.

An SQL implementation may detect the loss of the current SQL-connection during execution of any SQL-statement. When such a connection failure is detected, an exception condition is raised: *transaction rollback — statement completion unknown*. This exception condition indicates that the results of the actions performed in the SQL-server on behalf of the statement are unknown to the SQL-agent.

Similarly, an SQL-implementation may detect the loss of the current SQL-connection during the execution of a <commit statement>. When such a connection failure is detected, an exception condition is raised: *connection exception — transaction resolution unknown*. This exception condition indicates that the SQL-implementation cannot verify whether the SQL-transaction was committed successfully, rolled back, or left active.

A user may initiate an SQL-connection between the SQL-client associated with the SQL-agent and a specific SQL-server by executing a <connect statement>. Otherwise, an SQL-connection between the SQL-client and an implementation-defined default SQL-server is initiated when an externally-invoked procedure is called and no SQL-connection is current. The SQL-connection associated with an implementation-defined default SQL-server is called the *default SQL-connection*. An SQL-connection is terminated either by executing a <disconnect statement>, or following the last call to an externally-invoked procedure within the last active SQL-client module, or by the last execution of a <direct SQL statement> through the direct invocation of SQL. The mechanism and rules by which an SQL-implementation determines whether a call to an externally-invoked procedure is the last call within the last active SQL-client module and the mechanism and rules by which an SQL-implementation determines whether a direct invocation of SQL is the last execution of a <direct SQL statement> are implementation-defined.

An SQL-implementation shall support at least one SQL-connection and may require that the SQL-server be identified at the binding time chosen by the SQL-implementation. If an SQL-implementation permits more than one concurrent SQL-connection, then the SQL-agent may connect to more than one SQL-server and select the SQL-server by executing a <set connection statement>.

4.37 SQL-sessions

4.37.1 General description of SQL-sessions

This Subclause is modified by Subclause 4.12.1, “General description of SQL-sessions”, in ISO/IEC 9075-4.

An *SQL-session* spans the execution of a sequence of consecutive SQL-statements invoked either by a single user from a single SQL-agent or by the direct invocation of SQL. At any one time during an SQL-session, exactly one of the SQL-statements in this sequence is being executed and is said to be an *executing statement*. In some cases, an executing statement *ES* causes a nested sequence of consecutive SQL-statements to be executed as a direct result of *ES*; during that time, exactly one of these is also an executing statement and it in turn might similarly involve execution of a further nested sequence, and so on, indefinitely. An executing statement *ES*

such that no statement is executing as a direct result of *ES* is called the *innermost executing statement* of the SQL-session.

An SQL-session is associated with an SQL-connection. The SQL-session associated with the default SQL-connection is called the *default SQL-session*. An SQL-session is either *current* or *dormant*. The *current SQL-session* is the SQL-session associated with the current SQL-connection. A *dormant SQL-session* is an SQL-session that is associated with a dormant SQL-connection.

Within an SQL-session, SQL-client module declared local temporary tables are effectively created by <temporary table declaration>s contained in <SQL-client module definition>s. SQL-client module declared local temporary tables are accessible only to invocations of <externally-invoked procedure>s in the SQL-client module in which they are declared. The definitions of SQL-client module declared local temporary tables persist until the end of the SQL-session.

Within an SQL-session, locators are effectively created when a host parameter, a host variable, or an SQL parameter of an external routine that is specified as a binary large object locator, a character large object locator, a user-defined type locator, an array locator, or a multiset locator is assigned a value of binary large object type, character large object type, user-defined type, array type, or multiset type, respectively. These locators are part of the SQL-session context. A locator may be either valid or invalid. All locators remaining valid at the end of an SQL-session are marked invalid on termination of that SQL-session. A host variable that is a locator may be *holdable* or *nonholdable*.

4.37.2 SQL-session identification

An SQL-session has a unique implementation-dependent SQL-session identifier. This SQL-session identifier is different from the SQL-session identifier of any other concurrent SQL-session. The SQL-session identifier is used to effectively define implementation-defined schemas that contain the instances of any global temporary tables, created local temporary tables, or declared local temporary tables within the SQL-session.

An SQL-session is started as a result of successful execution of a <connect statement>, which sets the initial SQL-session user identifier to the value of the implicit or explicit <connection user name> contained in the <connect statement>.

An SQL-session initially has no SQL-session role name.

An SQL-session has an original time zone displacement and a current default time zone displacement, which are values of data type INTERVAL HOUR TO MINUTE. Both the original time zone displacement and the current default time zone displacement are initially set to the same implementation-defined value. The current default time zone displacement can subsequently be changed by successful execution of a <set local time zone statement>. The original time zone displacement cannot be changed. It is also possible to set the current default time zone displacement to equal the value of the original time zone displacement.

An SQL-session has a default catalog name that is used to effectively qualify unqualified <schema name>s that are contained in <preparable statement>s when those statements are prepared in the current SQL-session by either an <execute immediate statement> or a <prepare statement> or are contained in <direct SQL statement>s when those statements are invoked directly. The default catalog name is initially set to an implementation-defined value but can subsequently be changed by the successful execution of a <set catalog statement> or <set schema statement>. The default catalog name of an SQL-session may be determined by using the <general value specification> CURRENT_CATALOG.

An SQL-session has a default unqualified schema name that is used to effectively qualify unqualified <schema qualified name>s that are contained in <preparable statement>s when those statements are prepared in the

current SQL-session by either an <execute immediate statement> or a <prepare statement> or are contained in <direct SQL statement>s when those statements are invoked directly. The default unqualified schema name is initially set to an implementation-defined value but can subsequently be changed by the successful execution of a <set schema statement>. The default unqualified schema name of an SQL-session may be determined by using the <general value specification> CURRENT_SCHEMA.

4.37.3 SQL-session properties

This Subclause is modified by Subclause 4.16.1, “SQL-session properties”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 4.8.1, “SQL-session properties”, in ISO/IEC 9075-14.

An SQL-session has an SQL-path that is used to effectively qualify unqualified <routine name>s that are immediately contained in <routine invocation>s that are contained in <preparable statement>s when those statements are prepared in the current SQL-session by either an <execute immediate statement> or a <prepare statement> or are contained in <direct SQL statement>s when those statements are invoked directly. The SQL-path is initially set to an implementation-defined value, but can subsequently be changed by the successful execution of a <set path statement>.

The text defining the SQL-path can be referenced by using the <general value specification> CURRENT_PATH.

An SQL-session has a default transform group name and one or more user-defined type name—transform group name pairs that are used to identify the group of transform functions for every user-defined type that is referenced in <preparable statement>s when those statements are prepared in the current SQL-session by either an <execute immediate statement> or a <prepare statement> or is referenced in <direct SQL statement>s when those statements are invoked directly. The transform group name for a given user-defined type name is initially set to an implementation-defined value but can subsequently be changed by the successful execution of a <set transform group statement>.

The text defining the transform group names associated with the SQL-session can be referenced using two mechanisms: the <general value specification> “CURRENT_TRANSFORM_GROUP_FOR_TYPE <path-resolved user-defined type name>”, which evaluates to the name of the transform group associated with the specified data type, and the <general value specification> “CURRENT_DEFAULT_TRANSFORM_GROUP”, which evaluates to the name of the transform group associated with all types that have no type-specific transform group specified for them.

An SQL-session has a default character set name that is used to identify the character set in which <preparable statement>s are represented when those statements are prepared in the current SQL-session by either an <execute immediate statement> or a <prepare statement> or in which <direct SQL statement>s are represented when those statements are invoked directly. The default character set name is initially set to an implementation-defined value but can subsequently be changed by the successful execution of a <set names statement>.

[14] For each character set known to the SQL-implementation, an SQL-session has at most one SQL-session collation for that character set, to be used when the rules of Subclause 9.13, “Collation determination”, are applied. There are no SQL-session collations at the start of an SQL-session. The SQL-session collation for a character set can be set or changed by the successful execution of a <set session collation statement>.

An SQL-invoked routine is *active* as soon as an SQL-statement executed by an SQL-agent causes invocation of an SQL-invoked routine and ceases to be active when execution of that invocation is complete.

At any time during an SQL-session, *containing SQL* is said to be *permitted* or *not permitted*. Similarly, *reading SQL-data* is said to be *permitted* or *not permitted* and *modifying SQL-data* is said to be *permitted* or *not permitted*.

An SQL-session has *enduring characteristics*. The enduring characteristics of an SQL-session are initially the same as the default values for the corresponding SQL-session characteristics. The enduring characteristics are changed by successful execution of a <set session characteristics statement> that specifies one or more enduring characteristics. Enduring characteristics that are not specified in a <set session characteristics statement> are not changed in any way by the successful execution of that statement.

An SQL-session has the following enduring characteristics:

— *enduring transaction characteristics*.

Each of the enduring characteristics can be altered at any time in an SQL-session by executing an appropriate <set session characteristics statement>.

[14] An SQL-session has a stack of contexts that is preserved when an SQL-session is made dormant and restored when the SQL-session is made active. Each context in the stack comprises:

- The SQL-session identifier.
- The authorization stack.
- The identities of all instances of temporary tables.
- The original time zone displacement.
- The current default time zone displacement.
- The current constraint mode for each integrity constraint.
- The current access mode.
- The cursor instance descriptor of all open cursors.
- The current isolation level.
- The current SQL diagnostics area stack and its contents, along with the current condition area limit.
- The value of all valid locators.
- The value of the SQL-path for the current SQL-session.
- *A statement execution context*.
- *A routine execution context*.
- *Zero or more trigger execution contexts*.
- All prepared statements prepared during the current SQL-session and not deallocated.
- The current default catalog name.
- The current default unqualified schema name.
- The current default character set name.
- For each character set known to the SQL-implementation, the SQL-session collation, if any.
- The text defining the SQL-path.
- The contents of all SQL dynamic descriptor areas.

- The text defining the default transform group name.
- The text defining the user-defined type name—transform group name pair for each user-defined type explicitly set by the user.
- Each currently available result set sequence *RSS*, along with the specific name of an SQL-invoked procedure *SIP* and the name of the invoker of *SIP* for the invocation causing *RSS* to be brought into existence.

NOTE 74 — Result set sequences are defined in Subclause 4.27.5, “Result sets returned by SQL-invoked procedures”.

NOTE 75 — The use of the word “current” in the preceding list implies the values that are current in the SQL-session that is to be made dormant, and not the values that will become current in the SQL-session that will become the active SQL-session.

4.37.4 Execution contexts

Execution contexts augment an SQL-session context to cater for certain special circumstances that might pertain from time to time during invocations of SQL-statements. An execution context is either a statement execution context, a trigger execution context, or a routine execution context. There is always a *statement execution context*, a *routine execution context*, and zero or more *trigger execution contexts*. For certain SQL-statements, the statement execution context is always *atomic*; for others, it is always or sometimes non-atomic. A routine execution context is either atomic or non-atomic. Every trigger execution context is atomic. Statement execution contexts are described in Subclause 4.33.5, “SQL-statement atomicity and statement execution contexts”, routine execution contexts in Subclause 4.37.5, “Routine execution context”, and trigger execution contexts in Subclause 4.38.2, “Trigger execution”.

4.37.5 Routine execution context

A routine execution context consists of:

- An indication as to whether or not an SQL-invoked routine is active.
- An SQL-data access indication, which identifies what SQL-statements, if any, are allowed during the execution of an SQL-invoked routine. The SQL-data access indication is one of the following: does not possibly contain SQL, possibly contains SQL, possibly reads SQL-data, or possibly modifies SQL-data.
- An identification of the SQL-invoked routine that is active.
- The routine SQL-path derived from the routine SQL-path if the SQL-invoked routine that is active is an SQL routine and from the external routine SQL-path if the SQL-invoked routine that is active is an external routine.

An SQL-invoked routine is active as soon as an SQL-statement executed by an SQL-agent causes invocation of an SQL-invoked routine and ceases to be active when execution of that invocation is complete.

When an SQL-agent causes the invocation of an SQL-invoked routine, a new context for the current SQL-session is created and the values of the current context are preserved. When the execution of that SQL-invoked routine completes, the original context of the current SQL-session is restored and some SQL-session characteristics are reset.

If the routine execution context of the SQL-session indicates that an SQL-invoked routine is active, then the routine SQL-path included in the routine execution context of the SQL-session is used to effectively qualify

unqualified <routine name>s that are immediately contained in <routine invocation>s that are contained in a <preparable statement> or in a <direct SQL statement>.

4.38 Triggers

4.38.1 General description of triggers

A trigger is a specification for a given action to take place every time a given operation takes place on a given object. The action, known as a *triggered action*, is an SQL-procedure statement or a list of such statements. The object is either a persistent base table or a viewed table, known as the *subject table* of the trigger. The operation, known as a *trigger event*, is either deletion, insertion, or replacement of a collection of rows.

The triggered action is specified to take place either immediately before the triggering event, instead of it, or immediately after it, according to its specified *trigger action time*, BEFORE, INSTEAD OF, or AFTER. The trigger is a *BEFORE trigger*, an *INSTEAD OF trigger*, or an *AFTER trigger*, according to its trigger action time.

A trigger is either a *delete trigger*, an *insert trigger*, or an *update trigger*, according to the nature of its trigger event.

A trigger that is both a delete trigger and an INSTEAD OF trigger is a *delete INSTEAD OF trigger*. A trigger that is both an insert trigger and an INSTEAD OF trigger is an *insert INSTEAD OF trigger*. A trigger that is both an update trigger and an INSTEAD OF trigger is an *update INSTEAD OF trigger*.

Every trigger event arises as a consequence of executing some SQL-data change statement. That consequence might be direct, as for example when the SQL-data change statement is an <insert statement> operating on a base table, or indirect, as for example in the following cases:

- The SQL-data change statement is a <merge statement>.
- The SQL-data change statement operates on the referenced table of some foreign key whose referential action is CASCADE, SET NULL, or SET DEFAULT.
- The SQL-data change statement operates on a viewed table.

A triggered action is permitted to include SQL-data change statements that give rise to trigger events.

A collection of rows being deleted, inserted or replaced is known as a *transition table*. For a delete trigger there is just one transition table, known as an *old transition table*. For an insert trigger there is just one transition table, known as a *new transition table*. For an update trigger there is both an old transition table (the rows being replaced) and a new transition table (the replacement rows), these two tables having the same cardinality.

A reference to “the transition table” of a trigger is ambiguous in the case of an update trigger but whenever such a reference appears in this International Standard it is immaterial to which of the two transition tables it applies.

The triggered action can be specified to take place either just once when the trigger event takes place, in which case the trigger is a *statement-level trigger*, or once for each row of the transition table when the trigger event takes place, in which case the trigger is a *row-level trigger*.

4.38 Triggers

If the triggered action is specified to take place before the event, the trigger is a row-level trigger, and there is a new transition table, then the action can include statements whose effect is to alter the effect of the impending operation.

Special variables make the data in the transition table(s) available to the triggered action. For a statement-level trigger the variable is one whose value is a transition table. For a row-level trigger, the variable is a range variable, known as a *transition variable*. A transition variable ranges over the rows of a transition table, each row giving rise to exactly one execution of the triggered action, with the row in question assigned to the transition variable. A transition variable is either an *old transition variable* or a *new transition variable*, depending on the transition table over whose rows it ranges.

When there are two transition tables, old and new, each row in the new transition table is one that is derived by an update operation applied to exactly one row in the old transition table. Thus there is a 1:1 correspondence between the rows of the two tables. However, this correspondence is visible only to a row-level trigger, each invocation of which is able to access both the old and new transition variables, the new transition variable representing the result of applying the update operation in question to the row in the old transition variable.

A trigger is defined by a <trigger definition>, specifying the name of the trigger, its subject table, its trigger event, its trigger action time, whether it is statement-level or row-level, names as required for referencing transition tables or variables, and the triggered action.

A triggered action is a <triggered SQL statement> that is executed (either once for each affected row, in the case of a row-level trigger, or once for the whole trigger event in the case of a statement-level trigger) immediately before, instead of, or immediately after the trigger event takes place. The execution of a triggered action might cause the triggering of further triggered actions. It does so if it entails execution of an SQL-procedure statement whose effect causes the trigger event of some trigger to take place.

A trigger is described by a trigger descriptor. A trigger descriptor includes:

- The name of the trigger.
- The name of the subject table.
- The trigger action time (BEFORE, INSTEAD OF, or AFTER).
- The trigger event (INSERT, DELETE, or UPDATE).
- Whether the trigger is a statement-level trigger or a row-level trigger.
- Any old transition variable name, new transition variable name, old transition table name, or new transition table name.
- The triggered action.
- The trigger column list (possibly empty) for the trigger event.
- The triggered action column set of the triggered action.
- The timestamp of creation of the trigger.

The *order of execution* of a set of triggers is ascending by value of their timestamp of creation in their descriptors, such that the oldest trigger executes first. If one or more triggers have the same timestamp value, then their relative order of execution is implementation-defined.

A triggered action is always executed under the authorization of the owner of the schema that includes the trigger.

4.38.2 Trigger execution

During the execution of an SQL-statement S , zero or more *trigger execution contexts* exist, no more than one of which is *active* at any one time. A trigger execution context TEC_i comes into existence, becomes the active one, ceases to be active, and is destroyed as and when required under the General Rules for S .

If, while TEC_i is active, the General Rules for S require some new trigger execution context TEC_j to come into existence, then TEC_j replaces TEC_i as the active trigger execution context. TEC_i becomes active again when TEC_j is destroyed.

Multiple trigger execution contexts exist when the General Rules for S specify the execution of another SQL-procedure statement T before the execution of S is complete, and the General Rules for T require a new trigger execution context to come into existence.

A trigger execution context consists of a set of *state changes*. Within a trigger execution context, each state change is uniquely identified by a trigger event, a subject table, and a *column list*. The trigger event can be DELETE, INSERT, or UPDATE.

A state change SC consists of:

- A set of transitions.
- A trigger event.
- A subject table.
- A column list.
- A set (initially empty) of statement-level triggers *considered as executed* for SC .
- A set of row-level triggers, each paired with the set of rows in SC for which it is considered as executed.

What constitutes a transition depends on the trigger event. If the trigger event is DELETE, a transition is a row in the old transition table. If the trigger event is INSERT, a transition is a row in the new transition table. If the trigger event is UPDATE, a transition is a row OR in the old transition table paired with a row NR in the new transition table, such that NR is the row derived by applying a specified update operation to OR . OR and NR are the *old row* and the *new row*, respectively, of the transition.

A statement-level trigger that is considered as executed for a state change SC (in a given trigger execution context) is not subsequently executed for SC .

If a row-level trigger RLT is considered as executed for some row R in SC , then RLT is not subsequently executed for R .

A consequence of the execution of an SQL-data change statement is called an *SQL-update operation* if and only if that consequence causes at least one transition to arise in some state change.

A (possibly empty) old transition table exists if the trigger event is UPDATE or DELETE. It consists of a copy of each row that is to be updated in or deleted from the subject table. A (possibly empty) new transition table exists if the trigger event is UPDATE or INSERT. It consists of a copy of each row that results from updating a row in the subject table or is to be inserted into the subject table.

A <triggered action> may refer to the old transition table only if an <old transition table name> is specified for it in the <trigger definition>, and to the new transition table only if a <new transition table name> is specified for it in the <trigger definition>.

4.38 Triggers

The <triggered action> of a row-level trigger may refer to a range variable ranging over the rows of the old transition table only if an <old transition variable name> is specified for it in the <trigger definition>. Similarly, the <triggered action> of a row-level trigger may refer to a range variable ranging over the rows of the new transition table only if a <new transition variable name> is specified for it in the <trigger definition>. The scope of a transition variable or transition table name is the <triggered action> of the <trigger definition> that specifies it, excluding any <SQL schema statement>s that are contained in that <triggered action>.

When execution of an SQL-data change statement causes a trigger execution context TEC_i to come into existence, the set of state changes SSC_i in TEC_i is empty. Let SC_{ij} be a state change in SSC_i . Let TE be the trigger event (DELETE, INSERT, or UPDATE) of SC_{ij} . Let ST be the subject table of SC_{ij} .

If TE is INSERT or DELETE, then let PSC be a set whose only element is the empty set.

If TE is UPDATE, then:

- Let CL be the list of columns being updated by SSC_i .
- Let OC be the set of column names identifying the columns in CL .
- Let PSC be the set consisting of the empty set and every subset of the set of column names of ST that has at least one column that is in OC .

Let $PSCN$ be the number of elements in PSC . A state change SC_{ij} , for j varying from 1 (one) to $PSCN$, identified by TE , ST , and the j -th element in PSC , is added to SSC_i , provided that SSC_i does not already contain a state change corresponding to SC_{ij} . Transitions are added to SC_{ij} as specified by the General Rules of [Subclause 15.7](#), “Effect of deleting rows from base tables”, [Subclause 15.10](#), “Effect of inserting tables into base tables”, and [Subclause 15.13](#), “Effect of replacing rows in base tables”.

When a state change SC_{ij} arises in SSC_i , one or more triggers are *activated by* SC_{ij} . A trigger TR is activated by SC_{ij} if and only if the subject table of TR is the subject table of SC_{ij} , the trigger event of TR is the trigger event of SC_{ij} , and the set of column names listed in the trigger column list of TR is equivalent to the set of column names listed in SC_{ij} .

NOTE 76 — The trigger column list is included in the descriptor of TR ; it is empty if the trigger event is DELETE or INSERT. The trigger column list is also empty if the trigger event is UPDATE, but the <trigger event> of the <trigger definition> that defined TR does not specify a <trigger column list>.

For each state change SC_{ij} in TEC_i , the BEFORE triggers activated by SC_{ij} are executed before any of their triggering events take effect. When those triggering events have taken effect, any AFTER triggers activated by the state changes of TEC_i are executed.

The <triggered action> contained in a <trigger definition> for a row-level trigger can refer to columns of old transition variables and new transition variables. Such references can be specified as <column reference>s, which can be <target specification>s and <simple target specification>s when they refer to columns of the new transition variable in the triggered action of a BEFORE row-level trigger.

NOTE 77 — By using such <column reference>s as <assignment target>s (see [ISO9075-4]), the triggered action of a BEFORE trigger is able to cause certain SQL-data change statements to have different effects from those specified in the statements.

When an execution of the <triggered SQL statement> TSS of a triggered action is not successful, then an exception condition is raised and the SQL-statement that caused TSS to be executed has no effect on SQL-data or schemas.

4.39 Client-server operation

This Subclause is modified by Subclause 4.7, “Client-server operation”, in ISO/IEC 9075-3.

When an SQL-agent is active, it is bound in some implementation-defined manner to a single SQL-client. That SQL-client processes the explicit or implicit <SQL connection statement> for the first call to an externally-invoked procedure by an SQL-agent. The SQL-client communicates with, either directly or possibly through other agents such as RDA, one or more SQL-servers. An SQL-session involves an SQL-agent, an SQL-client, and a single SQL-server.

SQL-client modules associated with the SQL-agent exist in the SQL-environment containing the SQL-client associated with the SQL-agent.

Called <externally-invoked procedure>s and <direct SQL statement>s containing an <SQL connection statement> or an <SQL diagnostics statement> are processed by the SQL-client. Following the successful execution of a <connect statement> or a <set connection statement>, the SQL-client modules associated with the SQL-agent are effectively materialized with an implementation-dependent <SQL-client module name> in the SQL-server. Other called <externally-invoked procedure>s and <direct SQL statement>s are processed by the SQL-server.

A call by the SQL-agent to an <externally-invoked procedure> whose <SQL procedure statement> simply contains an <SQL diagnostics statement> fetches information from the specified diagnostics area in the diagnostics area stack associated with the SQL-client. Following the execution of an <SQL procedure statement> by an SQL-server, diagnostic information is passed in an implementation-dependent manner into the SQL-agent's diagnostics area stack in the SQL-client. The effect on diagnostic information of incompatibilities between the character repertoires supported by the SQL-client and SQL-server is implementation-dependent.

(Blank page)

5 Lexical elements

5.1 <SQL terminal character>

This Subclause is modified by Subclause 5.1, “<SQL terminal character>”, in ISO/IEC 9075-10.

Function

Define the terminal symbols of the SQL language and the elements of strings.

Format

```
<SQL terminal character> ::=
    <SQL language character>

<SQL language character> ::=
    <simple Latin letter>
  | <digit>
  | <SQL special character>

<simple Latin letter> ::=
    <simple Latin upper case letter>
  | <simple Latin lower case letter>

<simple Latin upper case letter> ::=
    A | B | C | D | E | F | G | H | I | J | K | L | M | N | O
  | P | Q | R | S | T | U | V | W | X | Y | Z

<simple Latin lower case letter> ::=
    a | b | c | d | e | f | g | h | i | j | k | l | m | n | o
  | p | q | r | s | t | u | v | w | x | y | z

<digit> ::=
    0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9

10 <SQL special character> ::=
    <space>
  | <double quote>
  | <percent>
  | <ampersand>
  | <quote>
  | <left paren>
  | <right paren>
  | <asterisk>
  | <plus sign>
  | <comma>
  | <minus sign>
  | <period>
```

5.1 <SQL terminal character>

	<solidus>
	<colon>
	<semicolon>
	<less than operator>
	<equals operator>
	<greater than operator>
	<question mark>
	<left bracket>
	<right bracket>
	<circumflex>
	<underscore>
	<vertical bar>
	<left brace>
	<right brace>

<space> ::=
!! See the Syntax Rules.

<double quote> ::=
"

<percent> ::=
%

<ampersand> ::=
&

<quote> ::=
,

<left paren> ::=
(

<right paren> ::=
)

<asterisk> ::=
*

<plus sign> ::=
+

<comma> ::=
,

<minus sign> ::=
-

<period> ::=
.

<solidus> ::=
/

<reverse solidus> ::=
\

<colon> ::=
:

```
<semicolon> ::=  
    ;  
  
<less than operator> ::=  
    <  
  
<equals operator> ::=  
    =  
  
<greater than operator> ::=  
    >  
  
<question mark> ::=  
    ?  
  
<left bracket or trigraph> ::=  
    <left bracket>  
    | <left bracket trigraph>  
  
<right bracket or trigraph> ::=  
    <right bracket>  
    | <right bracket trigraph>  
  
<left bracket> ::=  
    [  
  
<left bracket trigraph> ::=  
    ??(  
  
<right bracket> ::=  
    ]  
  
<right bracket trigraph> ::=  
    ??)  
  
<circumflex> ::=  
    ^  
  
<underscore> ::=  
    _  
  
<vertical bar> ::=  
    |  
  
<left brace> ::=  
    {  
  
<right brace> ::=  
    }
```

Syntax Rules

- 1) Every character set shall contain a <space> character that is equivalent to U+0020.

Access Rules

None.

General Rules

- 1) There is a one-to-one correspondence between the symbols contained in <simple Latin upper case letter> and the symbols contained in <simple Latin lower case letter> such that, for all i , the symbol defined as the i -th alternative for <simple Latin upper case letter> corresponds to the symbol defined as the i -th alternative for <simple Latin lower case letter>.

Conformance Rules

None.

5.2 <token> and <separator>

*This Subclause is modified by Subclause 5.1, “<token> and <separator>”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 5.1, “<token> and <separator>”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 5.2, “<token> and <separator>”, in ISO/IEC 9075-10.
This Subclause is modified by Subclause 5.1, “<token> and <separator>”, in ISO/IEC 9075-13.
This Subclause is modified by Subclause 5.1, “<token> and <separator>”, in ISO/IEC 9075-14.*

Function

Specify lexical units (tokens and separators) that participate in SQL language.

Format

```
<token> ::=
    <nondelimiter token>
  | <delimiter token>

<nondelimiter token> ::=
    <regular identifier>
  | <key word>
  | <unsigned numeric literal>
  | <national character string literal>
  | <binary string literal>
  | <large object length token>
  | <Unicode delimited identifier>
  | <Unicode character string literal>
  | <SQL language identifier>

<regular identifier> ::=
    <identifier body>

<identifier body> ::=
    <identifier start> [ <identifier part>... ]

<identifier part> ::=
    <identifier start>
  | <identifier extend>

<identifier start> ::=
    !! See the Syntax Rules.

<identifier extend> ::=
    !! See the Syntax Rules.

<large object length token> ::=
    <digit>... <multiplier>

<multiplier> ::=
    K
  | M
  | G
  | T
  | P
```

5.2 <token> and <separator>

```

<delimited identifier> ::=
    <double quote> <delimited identifier body> <double quote>

<delimited identifier body> ::=
    <delimited identifier part>...

<delimited identifier part> ::=
    <nondoublequote character>
  | <doublequote symbol>

<Unicode delimited identifier> ::=
    U<ampersand><double quote> <Unicode delimiter body> <double quote>
    <Unicode escape specifier>

<Unicode escape specifier> ::=
    [ UESCAPE <quote><Unicode escape character><quote> ]

<Unicode delimiter body> ::=
    <Unicode identifier part>...

<Unicode identifier part> ::=
    <delimited identifier part>
  | <Unicode escape value>

<Unicode escape value> ::=
    <Unicode 4 digit escape value>
  | <Unicode 6 digit escape value>
  | <Unicode character escape value>

<Unicode 4 digit escape value> ::=
    <Unicode escape character><hexit><hexit><hexit><hexit>

<Unicode 6 digit escape value> ::=
    <Unicode escape character><plus sign>
    <hexit><hexit><hexit><hexit><hexit><hexit>

<Unicode character escape value> ::=
    <Unicode escape character><Unicode escape character>

<Unicode escape character> ::=
    !! See the Syntax Rules.

<nondoublequote character> ::=
    !! See the Syntax Rules.

<doublequote symbol> ::=
    " !! two consecutive double quote characters

<delimiter token> ::=
    <character string literal>
  | <date string>
  | <time string>
  | <timestamp string>
  | <interval string>
  | <delimited identifier>
  | <SQL special character>
  | <not equals operator>
  | <greater than or equals operator>
  | <less than or equals operator>
  | <concatenation operator>

```

```
| <right arrow>
| <left bracket trigraph>
| <right bracket trigraph>
| <double colon>
| <double period>

<not equals operator> ::=
  <>

<greater than or equals operator> ::=
  >=

<less than or equals operator> ::=
  <=

<concatenation operator> ::=
  ||

<right arrow> ::=
  ->

<double colon> ::=
  ::

<double period> ::=
  ..

<separator> ::=
  { <comment> | <white space> }...

<white space> ::=
  !! See the Syntax Rules.

10 <comment> ::=
  <simple comment>
  | <bracketed comment>

<simple comment> ::=
  <simple comment introducer> [ <comment character>... ] <newline>

<simple comment introducer> ::=
  <minus sign><minus sign>

<bracketed comment> ::=
  <bracketed comment introducer>
  <bracketed comment contents>
  <bracketed comment terminator>

<bracketed comment introducer> ::=
  /*

<bracketed comment terminator> ::=
  */

<bracketed comment contents> ::=
  [ { <comment character> | <separator> }... ] !! See the Syntax Rules.

<comment character> ::=
  <nonquote character>
  | <quote>
```

5.2 <token> and <separator>

<newline> ::=

!! See the Syntax Rules.

<key word> ::=

<reserved word>

| <non-reserved word>

04 09 13 14 <non-reserved word> ::=

A | ABSOLUTE | ACTION | ADA | ADD | ADMIN | AFTER | ALWAYS | ASC
| ASSERTION | ASSIGNMENT | ATTRIBUTE | ATTRIBUTES

| BEFORE | BERNOULLI | BREADTH

| C | CASCADE | CATALOG | CATALOG_NAME | CHAIN | CHARACTER_SET_CATALOG
| CHARACTER_SET_NAME | CHARACTER_SET_SCHEMA | CHARACTERISTICS | CHARACTERS
| CLASS_ORIGIN | COBOL | COLLATION | COLLATION_CATALOG | COLLATION_NAME | COLLATION_SCHEMA
| COLUMN_NAME | COMMAND_FUNCTION | COMMAND_FUNCTION_CODE | COMMITTED
| CONDITION_NUMBER | CONNECTION | CONNECTION_NAME | CONSTRAINT_CATALOG | CONSTRAINT_NAME
| CONSTRAINT_SCHEMA | CONSTRAINTS | CONSTRUCTOR | CONTAINS | CONTINUE | CURSOR_NAME

| DATA | DATETIME_INTERVAL_CODE | DATETIME_INTERVAL_PRECISION | DEFAULTS | DEFERRABLE
| DEFERRED | DEFINED | DEFINER | DEGREE | DEPTH | DERIVED | DESC | DESCRIPTOR
| DIAGNOSTICS | DISPATCH | DOMAIN | DYNAMIC_FUNCTION | DYNAMIC_FUNCTION_CODE

| EQUALS | EXCLUDE | EXCLUDING

| FINAL | FIRST | FLAG | FOLLOWING | FORTRAN | FOUND

| G | GENERAL | GENERATED | GO | GOTO | GRANTED

| HIERARCHY

| IMMEDIATE | IMPLEMENTATION | INCLUDING | INCREMENT | INITIALLY | INPUT | INSTANCE
| INSTANTIABLE | INSTEAD | INVOKER | ISOLATION

| K | KEY | KEY_MEMBER | KEY_TYPE

| LAST | LENGTH | LEVEL | LOCATOR

| M | MAP | MATCHED | MAXVALUE | MESSAGE_LENGTH | MESSAGE_OCTET_LENGTH
| MESSAGE_TEXT | MINVALUE | MORE | MUMPS

| NAME | NAMES | NESTING | NEXT | NFC | NFD | NFKC | NFKD
| NORMALIZED | NULLABLE | NULLS | NUMBER

| OBJECT | OCTETS | OPTION | OPTIONS | ORDERING | ORDINALITY | OTHERS
| OUTPUT | OVERRIDING

| P | PAD | PARAMETER_MODE | PARAMETER_NAME | PARAMETER_ORDINAL_POSITION
| PARAMETER_SPECIFIC_CATALOG | PARAMETER_SPECIFIC_NAME | PARAMETER_SPECIFIC_SCHEMA
| PARTIAL | PASCAL | PATH | PLACING | PLI | PRECEDING | PRESERVE | PRIOR
| PRIVILEGES | PUBLIC

| READ | RELATIVE | REPEATABLE | RESTART | RESTRICT | RETURNED_CARDINALITY
| RETURNED_LENGTH | RETURNED_OCTET_LENGTH | RETURNED_SQLSTATE | ROLE
| ROUTINE | ROUTINE_CATALOG | ROUTINE_NAME | ROUTINE_SCHEMA | ROW_COUNT

| SCALE | SCHEMA | SCHEMA_NAME | SCOPE_CATALOG | SCOPE_NAME | SCOPE_SCHEMA
| SECTION | SECURITY | SELF | SEQUENCE | SERIALIZABLE | SERVER_NAME | SESSION
| SETS | SIMPLE | SIZE | SOURCE | SPACE | SPECIFIC_NAME | STATE | STATEMENT

```

| STRUCTURE | STYLE | SUBCLASS_ORIGIN

| T | TABLE_NAME | TEMPORARY | TIES | TOP_LEVEL_COUNT | TRANSACTION
| TRANSACTION_ACTIVE | TRANSACTIONS_COMMITTED | TRANSACTIONS_ROLLED_BACK
| TRANSFORM | TRANSFORMS | TRIGGER_CATALOG | TRIGGER_NAME | TRIGGER_SCHEMA | TYPE

| UNBOUNDED | UNCOMMITTED | UNDER | UNNAMED | USAGE | USER_DEFINED_TYPE_CATALOG
| USER_DEFINED_TYPE_CODE | USER_DEFINED_TYPE_NAME | USER_DEFINED_TYPE_SCHEMA

| VIEW

| WORK | WRITE

| ZONE

04 09 13 14 <reserved word> ::=
ABS | ALL | ALLOCATE | ALTER | AND | ANY | ARE | ARRAY | ARRAY_AGG | AS
| ASENSITIVE | ASYMMETRIC | AT | ATOMIC | AUTHORIZATION | AVG

| BEGIN | BETWEEN | BIGINT | BINARY | BLOB | BOOLEAN | BOTH | BY

| CALL | CALLED | CARDINALITY | CASCADED | CASE | CAST | CEIL | CEILING
| CHAR | CHAR_LENGTH | CHARACTER | CHARACTER_LENGTH | CHECK | CLOB | CLOSE
| COALESCE | COLLATE | COLLECT | COLUMN | COMMIT | CONDITION | CONNECT
| CONSTRAINT | CONVERT | CORR | CORRESPONDING | COUNT | COVAR_POP | COVAR_SAMP
| CREATE | CROSS | CUBE | CUME_DIST | CURRENT | CURRENT_CATALOG | CURRENT_DATE
| CURRENT_DEFAULT_TRANSFORM_GROUP | CURRENT_PATH | CURRENT_ROLE | CURRENT_SCHEMA
| CURRENT_TIME | CURRENT_TIMESTAMP | CURRENT_TRANSFORM_GROUP_FOR_TYPE
| CURRENT_USER | CURSOR | CYCLE

| DATE | DAY | DEALLOCATE | DEC | DECIMAL | DECLARE | DEFAULT | DELETE
| DENSE_RANK | Deref | DESCRIBE | DETERMINISTIC | DISCONNECT | DISTINCT
| DOUBLE | DROP | DYNAMIC

| EACH | ELEMENT | ELSE | END | END-EXEC | ESCAPE | EVERY | EXCEPT | EXEC
| EXECUTE | EXISTS | EXP | EXTERNAL | EXTRACT

| FALSE | FETCH | FILTER | FLOAT | FLOOR | FOR | FOREIGN | FREE | FROM
| FULL | FUNCTION | FUSION

| GET | GLOBAL | GRANT | GROUP | GROUPING

| HAVING | HOLD | HOUR

| IDENTITY | IN | INDICATOR | INNER | INOUT | INSENSITIVE | INSERT
| INT | INTEGER | INTERSECT | INTERSECTION | INTERVAL | INTO | IS

| JOIN

| LANGUAGE | LARGE | LATERAL | LEADING | LEFT | LIKE | LIKE_REGEX
| LN | LOCAL | LOCALTIME | LOCALTIMESTAMP | LOWER

| MATCH | MAX | MEMBER | MERGE | METHOD | MIN | MINUTE
| MOD | MODIFIES | MODULE | MONTH | MULTISSET

| NATIONAL | NATURAL | NCHAR | NCLOB | NEW | NO | NONE | NORMALIZE | NOT
| NULL | NULLIF | NUMERIC

| OCTET_LENGTH | OCCURRENCES_REGEX | OF | OLD | ON | ONLY | OPEN | OR
| ORDER | OUT | OUTER | OVER | OVERLAPS | OVERLAY

```

	PARAMETER		PARTITION		PERCENT_RANK		PERCENTILE_CONT		PERCENTILE_DISC									
	POSITION		POSITION_REGEX		POWER		PRECISION		PREPARE		PRIMARY							
	PROCEDURE																	
	RANGE		RANK		READS		REAL		RECURSIVE		REF		REFERENCES		REFERENCING			
	REGR_AVGX		REGR_AVGY		REGR_COUNT		REGR_INTERCEPT		REGR_R2		REGR_SLOPE							
	REGR_SXX		REGR_SXY		REGR_SYY		RELEASE		RESULT		RETURN		RETURNS					
	REVOKE		RIGHT		ROLLBACK		ROLLUP		ROW		ROW_NUMBER		ROWS					
	SAVEPOINT		SCOPE		SCROLL		SEARCH		SECOND		SELECT		SENSITIVE					
	SESSION_USER		SET		SIMILAR		SMALLINT		SOME		SPECIFIC		SPECIFICTYPE					
	SQL		SQLEXCEPTION		SQLSTATE		SQLWARNING		SQRT		START		STATIC					
	STDDEV_POP		STDDEV_SAMP		SUBMULTISET		SUBSTRING		SUBSTRING_REGEX		SUM							
	SYMMETRIC		SYSTEM		SYSTEM_USER													
	TABLE		TABLESAMPLE		THEN		TIME		TIMESTAMP		TIMEZONE_HOUR		TIMEZONE_MINUTE					
	TO		TRAILING		TRANSLATE		TRANSLATE_REGEX		TRANSLATION		TREAT							
	TRIGGER		TRUNCATE		TRIM		TRUE											
	UESCAPE		UNION		UNIQUE		UNKNOWN		UNNEST		UPDATE		UPPER		USER		USING	
	VALUE		VALUES		VAR_POP		VAR_SAMP		VARBINARY		VARCHAR		VARYING					
	WHEN		WHENEVER		WHERE		WIDTH_BUCKET		WINDOW		WITH		WITHIN		WITHOUT			
	YEAR																	

Syntax Rules

- 1) An <identifier start> is any character in the Unicode General Category classes “Lu”, “Ll”, “Lt”, “Lm”, “Lo”, or “Nl”.

NOTE 78 — The Unicode General Category classes “Lu”, “Ll”, “Lt”, “Lm”, “Lo”, and “Nl” are assigned to Unicode characters that are, respectively, upper-case letters, lower-case letters, title-case letters, modifier letters, other letters, and letter numbers.

- 2) An <identifier extend> is U+00B7, “Middle Dot”, or any character in the Unicode General Category classes “Mn”, “Mc”, “Nd”, “Pc”, or “Cf”.

NOTE 79 — The Unicode General Category classes “Mn”, “Mc”, “Nd”, “Pc”, and “Cf” are assigned to Unicode characters that are, respectively, nonspacing marks, spacing combining marks, decimal numbers, connector punctuations, and formatting codes.

- 3) <white space> is any consecutive sequence of characters each of which satisfies the definition of white space found in [Subclause 3.1.6, “Definitions provided in Part 2”](#).

- 4) <newline> is the implementation-defined end-of-line indicator.

NOTE 80 — <newline> is typically represented by U+000A (“Line Feed”) and/or U+000D (“Carriage Return”); however, this representation is not required by ISO/IEC 9075.

- 5) With the exception of the <space> character explicitly contained in <binary string literal>, <timestamp string>, and <interval string>, a <token>, other than a <character string literal>, a <national character string literal>, a <Unicode character string literal>, a <delimited identifier>, or a <Unicode delimited identifier> shall not contain a <space> character or other <separator>.

- 6) A <nondoublequote character> is any character of the source language character set other than a <double quote>.

NOTE 81 — “source language character set” is defined in [Subclause 4.8.1, “Host languages”](#), in [ISO9075-1].

- 7) Any <token> may be followed by a <separator>. A <nondelimiter token> shall be followed by a <delimiter token> or a <separator>.

NOTE 82 — If the Format does not allow a <nondelimiter token> to be followed by a <delimiter token>, then that <nondelimiter token> shall be followed by a <separator>.

- 8) There shall be no <separator> separating the <minus sign>s of a <simple comment introducer>.
- 9) There shall be no <separator> separating any two <digit>s or separating a <digit> and <multiplier> of a <large object length token>.
- 10) Within a <bracketed comment contents>, any <solidus> immediately followed by an <asterisk> without any intervening <separator> shall be considered to be the <bracketed comment introducer> of a <separator> that is a <bracketed comment>.

NOTE 83 — Conforming programs should not place <simple comment> within a <bracketed comment> because if such a <simple comment> contains the sequence of characters “*/” without a preceding “/*” in the same <simple comment>, it will prematurely terminate the containing <bracketed comment>.

- 11) SQL text containing one or more instances of <comment> is equivalent to the same SQL text with the <comment> replaced with <newline>.
- 12) In a <regular identifier>, the number of <identifier part>s shall be less than 128.
- 13) The <delimited identifier body> of a <delimited identifier> shall not comprise more than 128 <delimited identifier part>s.
- 14) In a <Unicode delimited identifier>, there shall be no <separator> between the 'U' and the <ampersand> nor between the <ampersand> and the <double quote>.
- 15) In a <Unicode delimited identifier>, the introductory 'U' may be represented either in upper case (as 'U') or in lower case (as 'u').
- 16) <Unicode escape character> shall be a single character from the source language character set other than a <hexit>, <plus sign>, <quote>, <double quote>, or <white space>.
- 17) If the source language character set contains <reverse solidus>, then let *DEC* be <reverse solidus>; otherwise, let *DEC* be an implementation-defined character from the source language character set that is not a <hexit>, <plus sign>, <quote>, <double quote>, or <white space>.
- 18) If a <Unicode escape specifier> does not contain <Unicode escape character>, then “UESCAPE <quote>*DEC*<quote>” is implicit.
- 19) In a <Unicode escape value> there shall be no <separator> between the <Unicode escape character> and the first <hexit>, nor between any of the <hexit>s.
- 20) The <Unicode delimiter body> of a <Unicode delimited identifier> shall not comprise more than 128 <Unicode identifier part>s.
- 21) <Unicode 4 digit escape value> '<Unicode escape character>xyzw' is equivalent to the character at the Unicode code point specified by U+xyzw.
- 22) <Unicode 6 digit escape value> '<Unicode escape character>+xyzwrs' is equivalent to the character at the Unicode code point specified by U+xyzwrs.

NOTE 84 — The 6-hexit notation is derived by taking the UCS-4 notation defined in [ISO10646] and removing the leading two hexits, whose values are always 0 (zero).

- 23) <Unicode character escape value> is equivalent to a single instance of <Unicode escape character>.

- 24) For every <identifier body> *IB* there is exactly one corresponding case-normal form *CNF*. *CNF* is an <identifier body> derived from *IB* as follows:

Let *n* be the number of characters in *IB*. For *i* ranging from 1 (one) to *n*, the *i*-th character *M_i* of *IB* is transliterated into the corresponding character or characters of *CNF* as follows:

Case:

- a) If *M_i* is a lower case character or a title case character for which an equivalent upper case sequence *U* is defined by Unicode, then let *j* be the number of characters in *U*; the next *j* characters of *CNF* are *U*.
- b) Otherwise, the next character of *CNF* is *M_i*.

- 25) The case-normal form of the <identifier body> of a <regular identifier> is used for purposes such as and including determination of identifier equivalence, representation in the Definition and Information Schemas, and representation in diagnostics areas.

NOTE 85 — The Information Schema and Definition Schema are defined in [ISO9075-11].

NOTE 86 — Any lower-case letters for which there are no upper-case equivalents are left in their lower-case form.

- 26) The case-normal form of <regular identifier> shall not be equal, according to the comparison rules in Subclause 8.2, “<comparison predicate>”, to any <reserved word> (with every letter that is a lower-case letter replaced by the corresponding upper-case letter or letters), treated as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER.
- 27) Two <regular identifier>s are equivalent if the case-normal forms of their <identifier body>s, considered as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER and an implementation-defined collation *IDC* that is sensitive to case, compare equally according to the comparison rules in Subclause 8.2, “<comparison predicate>”.
- 28) A <regular identifier> and a <delimited identifier> are equivalent if the case-normal form of the <identifier body> of the <regular identifier> and the <delimited identifier body> of the <delimited identifier> (with all occurrences of <quote> replaced by <quote symbol> and all occurrences of <doublequote symbol> replaced by <double quote>), considered as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER and *IDC*, compare equally according to the comparison rules in Subclause 8.2, “<comparison predicate>”.
- 29) Two <delimited identifier>s are equivalent if their <delimited identifier body>s, considered as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER and an implementation-defined collation that is sensitive to case, compare equally according to the comparison rules in Subclause 8.2, “<comparison predicate>”.
- 30) Two <Unicode delimited identifier>s are equivalent if their <Unicode delimiter body>s, considered as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER and an implementation-defined collation that is sensitive to case, compare equally according to the comparison rules in Subclause 8.2, “<comparison predicate>”.
- 31) A <Unicode delimited identifier> and a <delimited identifier> are equivalent if their <Unicode delimiter body> and <delimited identifier body>, respectively, each considered as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER and an implementation-defined collation that is sensitive to case, compare equally according to the comparison rules in Subclause 8.2, “<comparison predicate>”.
- 32) For the purposes of identifying <key word>s, any <simple Latin lower case letter> contained in a candidate <key word> shall be effectively treated as the corresponding <simple Latin upper case letter>.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature F391, “Long identifiers”, in a <regular identifier>, the number of <identifier part>s shall be less than 18.
- 2) Without Feature F391, “Long identifiers”, the <delimited identifier body> of a <delimited identifier> shall not comprise more than 18 <delimited identifier part>s.

NOTE 87 — Not every character set supported by a conforming SQL-implementation necessarily contains every character associated with <identifier start> and <identifier part> that is identified in the Syntax Rules of this Subclause. No conforming SQL-implementation shall be required to support in <identifier start> or <identifier part> any character identified in the Syntax Rules of this Subclause unless that character belongs to the character set in use for an SQL-client module or in SQL-data.

- 3) Without Feature T351, “Bracketed comments”, conforming SQL language shall not contain a <bracketed comment>.
- 4) Without Feature F392, “Unicode escapes in identifiers”, conforming SQL language shall not contain a <Unicode delimited identifier>.
- 5) Without Feature T043, “Multiplier T”, in conforming SQL language, a <multiplier> shall not be T.
- 6) Without Feature T044, “Multiplier P”, in conforming SQL language, a <multiplier> shall not be P.

5.3 <literal>

Function

Specify a non-null value.

Format

```

<literal> ::=
    <signed numeric literal>
  | <general literal>

<unsigned literal> ::=
    <unsigned numeric literal>
  | <general literal>

<general literal> ::=
    <character string literal>
  | <national character string literal>
  | <Unicode character string literal>
  | <binary string literal>
  | <datetime literal>
  | <interval literal>
  | <boolean literal>

<character string literal> ::=
    [ <introducer><character set specification> ]
    <quote> [ <character representation>... ] <quote>
    [ { <separator> <quote> [ <character representation>... ] <quote> }... ]

<introducer> ::=
    <underscore>

<character representation> ::=
    <nonquote character>
  | <quote symbol>

<nonquote character> ::=
    !! See the Syntax Rules.

<quote symbol> ::=
    <quote><quote>

<national character string literal> ::=
    N <quote> [ <character representation>... ]
    <quote> [ { <separator> <quote> [ <character representation>... ] <quote> }... ]

<Unicode character string literal> ::=
    [ <introducer><character set specification> ]
    U<ampersand><quote> [ <Unicode representation>... ] <quote>
    [ { <separator> <quote> [ <Unicode representation>... ] <quote> }... ]
    <Unicode escape specifier>

<Unicode representation> ::=
    <character representation>
  | <Unicode escape value>

```

```

<binary string literal> ::=
  X <quote> [ <space>... ] [ { <hexit> [ <space>... ] <hexit> [ <space>... ] }... ] <quote>
    [ { <separator> <quote> [ <space>... ] [ { <hexit> [ <space>... ]
      <hexit> [ <space>... ] }... ] <quote> }... ]

<hexit> ::=
  <digit> | A | B | C | D | E | F | a | b | c | d | e | f

<signed numeric literal> ::=
  [ <sign> ] <unsigned numeric literal>

<unsigned numeric literal> ::=
  <exact numeric literal>
  | <approximate numeric literal>

<exact numeric literal> ::=
  <unsigned integer> [ <period> [ <unsigned integer> ] ]
  | <period> <unsigned integer>

<sign> ::=
  <plus sign>
  | <minus sign>

<approximate numeric literal> ::=
  <mantissa> E <exponent>

<mantissa> ::=
  <exact numeric literal>

<exponent> ::=
  <signed integer>

<signed integer> ::=
  [ <sign> ] <unsigned integer>

<unsigned integer> ::=
  <digit>...

<datetime literal> ::=
  <date literal>
  | <time literal>
  | <timestamp literal>

<date literal> ::=
  DATE <date string>

<time literal> ::=
  TIME <time string>

<timestamp literal> ::=
  TIMESTAMP <timestamp string>

<date string> ::=
  <quote> <unquoted date string> <quote>

<time string> ::=
  <quote> <unquoted time string> <quote>

<timestamp string> ::=
  <quote> <unquoted timestamp string> <quote>

```

5.3 <literal>

```

<time zone interval> ::=
    <sign> <hours value> <colon> <minutes value>

<date value> ::=
    <years value> <minus sign> <months value> <minus sign> <days value>

<time value> ::=
    <hours value> <colon> <minutes value> <colon> <seconds value>

<interval literal> ::=
    INTERVAL [ <sign> ] <interval string> <interval qualifier>

<interval string> ::=
    <quote> <unquoted interval string> <quote>

<unquoted date string> ::=
    <date value>

<unquoted time string> ::=
    <time value> [ <time zone interval> ]

<unquoted timestamp string> ::=
    <unquoted date string> <space> <unquoted time string>

<unquoted interval string> ::=
    [ <sign> ] { <year-month literal> | <day-time literal> }

<year-month literal> ::=
    <years value> [ <minus sign> <months value> ]
    | <months value>

<day-time literal> ::=
    <day-time interval>
    | <time interval>

<day-time interval> ::=
    <days value> [ <space> <hours value> [ <colon> <minutes value>
        [ <colon> <seconds value> ] ] ]

<time interval> ::=
    <hours value> [ <colon> <minutes value> [ <colon> <seconds value> ] ]
    | <minutes value> [ <colon> <seconds value> ]
    | <seconds value>

<years value> ::=
    <datetime value>

<months value> ::=
    <datetime value>

<days value> ::=
    <datetime value>

<hours value> ::=
    <datetime value>

<minutes value> ::=
    <datetime value>

<seconds value> ::=

```

```
<seconds integer value> [ <period> [ <seconds fraction> ] ]  
  
<seconds integer value> ::=  
    <unsigned integer>  
  
<seconds fraction> ::=  
    <unsigned integer>  
  
<datetime value> ::=  
    <unsigned integer>  
  
<boolean literal> ::=  
    TRUE  
    | FALSE  
    | UNKNOWN
```

Syntax Rules

- 1) In a <character string literal> or <national character string literal>, the sequence:

```
<quote> <character representation>... <quote><separator>  
<quote> <character representation>... <quote>
```

is equivalent to the sequence

```
<quote> <character representation>... <character representation>... <quote>
```

NOTE 88 — The <character representation>s in the equivalent sequence are in the same sequence and relative sequence as in the original <character string literal>.

- 2) In a <Unicode character string literal>, the sequence:

```
<quote> <Unicode representation>... <quote>  
<separator> <quote> <Unicode representation>... <quote>
```

is equivalent to the sequence:

```
<quote> <Unicode representation>... <Unicode representation>... <quote>
```

- 3) In a <Unicode character string literal>, the introductory 'U' may be represented either in upper case (as 'U') or in lower case (as 'u').

- 4) In a <binary string literal>, the sequence

```
<quote> [ <space>... ] { <hexit> [ <space>... ]  
<hexit> [ <space>... ] }... <quote>
```

is equivalent to the sequence

```
<quote> { <hexit> <hexit> }... <quote>
```

NOTE 89 — The <hexit>s in the equivalent sequence are in the same sequence and relative sequence as in the original <binary string literal>.

- 5) In a <binary string literal>, the sequence

5.3 <literal>

```
<quote> { <hexit> <hexit> }... <quote> <separator>
<quote> { <hexit> <hexit> }... <quote>
```

is equivalent to the sequence

```
<quote> { <hexit> <hexit> }... { <hexit> <hexit> }... <quote>
```

NOTE 90 — The <hexit>s in the equivalent sequence are in the same sequence and relative sequence as in the original <binary string literal>.

- 6) In a <binary string literal>, the introductory 'X' may be represented either in upper case (as 'X') or in lower case (as 'x').
 - 7) In a <character string literal>, <national character string literal>, <Unicode character string literal>, or <binary string literal>, a <separator> shall contain a <newline>.
 - 8) A <national character string literal> is equivalent to a <character string literal> with the “N” replaced by “<introducer><character set specification>”, where “<character set specification>” is an implementation-defined <character set name>.
 - 9) In a <national character string literal>, the introductory 'N' may be represented either in upper case (as 'N') or in lower case (as 'n').
 - 10) In a <Unicode character string literal> that specifies “<introducer><character set specification>”, there shall be no <separator> between the <introducer> and the <character set specification>.
 - 11) In a <Unicode character string literal>, there shall be no <separator> between the “U” and the <ampersand> nor between the <ampersand> and the <quote>.
 - 12) The character set of a <Unicode character string literal> that specifies “<introducer><character set specification>” is the character set specified by the <character set specification>. The character set of a <Unicode character string literal> that does not specify “<introducer><character set specification>” is the character set of the SQL-client module that contains the <Unicode character string literal>.
 - 13) A <Unicode character string literal> is equivalent to a <character string literal> in which every <Unicode escape value> has been replaced with the equivalent Unicode character. The set of characters contained in the <Unicode character string literal> shall be wholly contained in the character set of the <Unicode character string literal>.
- NOTE 91 — The requirement for “wholly contained” applies after the replacement of <Unicode escape value>s with equivalent Unicode characters.
- 14) Each <character representation> is a character of the source language character set. The value of a <character string literal>, viewed as a string in the source language character set, shall be equivalent to a character string of the implicit or explicit character set of the <character string literal> or <national character string literal>.
- NOTE 92 — “source language character set” is defined in Subclause 4.8.1, “Host languages”, in [ISO9075-1].
- 15) A <nonquote character> is one of:
 - a) Any character of the source language character set other than a <quote>.
 - b) Any character other than a <quote> in the character set identified by the <character set specification> or implied by “N”.
 - 16) Case:

- a) If a <character set specification> is not specified in a <character string literal>, then the set of characters contained in the <character string literal> shall be wholly contained in the character set of the <SQL-client module definition> that contains the <character string literal>.
 - b) Otherwise, there shall be no <separator> between the <introducer> and the <character set specification>, and the set of characters contained in the <character string literal> shall be wholly contained in the character set specified by the <character set specification>.
- 17) The declared type of a <character string literal> is fixed-length character string. The length of a <character string literal> is the number of <character representation>s that it contains. Each <quote symbol> contained in <character string literal> represents a single <quote> in both the value and the length of the <character string literal>. The two <quote>s contained in a <quote symbol> shall not be separated by any <separator>.

NOTE 93 — <character string literal>s are allowed to be zero-length strings (*i.e.*, to contain no characters) even though it is not permitted to declare a <data type> that is CHARACTER with <character length> 0 (zero).

- 18) The character set of a <character string literal> is

Case:

- a) If the <character string literal> specifies a <character set specification>, then the character set specified by that <character set specification>.
 - b) Otherwise, the character set of the SQL-client module that contains the <character string literal>.
- 19) The declared type collation of a <character string literal> is the character set collation, and the collation derivation is *implicit*.
- 20) It is implementation-defined whether the declared type of a <binary string literal> is a fixed-length binary string type, a variable-length binary string type, or a binary large object string type. Each <hexit> appearing in the literal is equivalent to a quartet of bits: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, and F are interpreted as 0000, 0001, 0010, 0011, 0100, 0101, 0110, 0111, 1000, 1001, 1010, 1011, 1100, 1101, 1110, and 1111, respectively. The <hexit>s a, b, c, d, e, and f have respectively the same values as the <hexit>s A, B, C, D, E, and F.
- 21) An <exact numeric literal> without a <period> has an implied <period> following the last <digit>.
- 22) The declared type of an <exact numeric literal> *ENL* is an implementation-defined exact numeric type whose scale is the number of <digit>s to the right of the <period>. There shall be an exact numeric type capable of representing the value of *ENL* exactly.
- 23) The declared type of an <approximate numeric literal> *ANL* is an implementation-defined approximate numeric type. The value of *ANL* shall not be greater than the maximum value nor less than the minimum value that can be represented by the approximate numeric types.
- NOTE 94 — Thus the only syntax error for an <approximate numeric literal> is what is commonly known as “overflow”; there is no syntax error for specifying more significant digits than the SQL-implementation can represent internally, nor for specifying a value that has no exact equivalent in the SQL-implementation's internal representation. (“Underflow”, *i.e.*, specifying a nonzero value so close to 0 (zero) that the closest representation in the SQL-implementation's internal representation is 0E0, is a special case of the latter condition, and is not a syntax error.)
- 24) In an <approximate numeric literal>, the exponent indicator 'E' may be represented either in upper case (as 'E') or in lower case (as 'e').
- 25) The declared type of a <date literal> is DATE.
- 26) The declared type of a <time literal> that does not specify <time zone interval> is TIME(*P*) WITHOUT TIME ZONE, where *P* is the number of digits in <seconds fraction>, if specified, and 0 (zero) otherwise.

5.3 <literal>

The declared type of a <time literal> that specifies <time zone interval> is TIME(*P*) WITH TIME ZONE, where *P* is the number of digits in <seconds fraction>, if specified, and 0 (zero) otherwise.

- 27) The declared type of a <timestamp literal> that does not specify <time zone interval> is TIMESTAMP(*P*) WITHOUT TIME ZONE, where *P* is the number of digits in <seconds fraction>, if specified, and 0 (zero) otherwise. The declared type of a <timestamp literal> that specifies <time zone interval> is TIMESTAMP(*P*) WITH TIME ZONE, where *P* is the number of digits in <seconds fraction>, if specified, and 0 (zero) otherwise.
- 28) If <time zone interval> is not specified, then the effective <time zone interval> of the datetime data type is the current default time zone displacement for the SQL-session.
- 29) Let *datetime component* be either <years value>, <months value>, <days value>, <hours value>, <minutes value>, or <seconds value>.
- 30) Let *N* be the number of <primary datetime field>s in the precision of the <interval literal>, as specified by <interval qualifier>.

The <interval literal> being defined shall contain *N* datetime components.

The declared type of <interval literal> specified with an <interval qualifier> is INTERVAL with the <interval qualifier>.

Each datetime component shall have the precision specified by the <interval qualifier>.

- 31) Within a <datetime literal>, the <years value> shall contain four digits. The <seconds integer value> and other datetime components, with the exception of <seconds fraction>, shall each contain two digits.
- 32) Within the definition of a <datetime literal>, the <datetime value>s are constrained by the natural rules for dates and times according to the Gregorian calendar.
- 33) Within the definition of an <interval literal>, the <datetime value>s are constrained by the natural rules for intervals according to the Gregorian calendar.
- 34) Within the definition of an <interval literal> that contains a <year-month literal>, the <interval qualifier> shall not specify DAY, HOUR, MINUTE, or SECOND. Within the definition of an <interval literal> that contains a <day-time literal>, the <interval qualifier> shall not specify YEAR or MONTH.
- 35) Within the definition of a <datetime literal>, the value of the <time zone interval> shall be in the range −14:00 to +14:00.

Access Rules

None.

General Rules

- 1) The value of a <character string literal> is the result of transliterating the sequence of <character representation>s that it contains from the source language character set to the implicit or explicit character set of the <character string literal>.
- 2) If the character repertoire of a <character string literal> *US* is UCS, then its value is replaced by NORMALIZE(*US*).

- 3) Except when it is contained in an <exact numeric literal>, the value of an <unsigned integer> is the numeric value determined by application of the normal mathematical interpretation of positional decimal notation to the string of <digit>s that constitutes the <unsigned integer>.
- 4) The value of an <exact numeric literal> is the numeric value determined by application of the normal mathematical interpretation of positional decimal notation to the source characters that constitute the <exact numeric literal>.
- 5) Let *ANL* be an <approximate numeric literal>. Let *ANDT* be the declared type of *ANL*. Let *ANV* be the product of the exact numeric value represented by the <mantissa> of *ANL* and the number obtained by raising the number 10 to the power of the exact numeric value represented by the <exponent> of *ANL*. If *ANV* is a member of *ANDT*, then the value of *ANL* is *ANV*; otherwise, the value of *ANL* is a member of *ANDT* obtained from *ANV* by rounding or truncation.
- 6) The <sign> in a <signed numeric literal> or an <interval literal> is a monadic arithmetic operator. The monadic arithmetic operators + and – specify monadic plus and monadic minus, respectively. If neither monadic plus nor monadic minus are specified in a <signed numeric literal> or an <interval literal>, or if monadic plus is specified, then the literal is positive. If monadic minus is specified in a <signed numeric literal> or <interval literal>, then the literal is negative. If <sign> is specified in both possible locations in an <interval literal>, then the sign of the literal is determined by normal mathematical interpretation of multiple sign operators.
- 7) Let *V* be the integer value of the <unsigned integer> contained in <seconds fraction> and let *N* be the number of digits in the <seconds fraction> respectively. The resultant value of the <seconds fraction> is effectively determined as follows.

Case:

- a) If <seconds fraction> is specified within the definition of a <datetime literal>, then the effective value of the <seconds fraction> is $V * 10^{-N}$ seconds.
- b) If <seconds fraction> is specified within the definition of an <interval literal>, then let *M* be the <interval fractional seconds precision> specified in the <interval qualifier>.

Case:

- i) If $N < M$, then let *V1* be $V * 10^{M-N}$; the effective value of the <seconds fraction> is $V1 * 10^{-M}$ seconds.
 - ii) If $N > M$, then let *V2* be the integer part of the quotient of $V/10^{N-M}$; the effective value of the <seconds fraction> is $V2 * 10^{-M}$ seconds.
 - iii) Otherwise, the effective value of the <seconds fraction> is $V * 10^{-M}$ seconds.
- 8) The *i*-th datetime component in a <datetime literal> or <interval literal> assigns the value of the datetime component to the *i*-th <primary datetime field> in the <datetime literal> or <interval literal>.
 - 9) If <time zone interval> is specified, then the time and timestamp values in <time literal> and <timestamp literal> represent a datetime in the specified time zone.
 - 10) If <date value> is specified, then it is interpreted as a date in the Gregorian calendar. If <time value> is specified, then it is interpreted as a time of day. Let *DV* be the value of the <datetime literal>, disregarding <time zone interval>.

Case:

5.3 <literal>

- a) If <time zone interval> is specified, then let *TZI* be the value of the interval denoted by <time zone interval>. The value of the <datetime literal> is $DV - TZI$, with time zone displacement *TZI*.
- b) Otherwise, the value of the <datetime literal> is *DV*.

NOTE 95 — If <time zone interval> is specified, then a <time literal> or <timestamp literal> is interpreted as local time with the specified time zone displacement. However, it is effectively converted to UTC while retaining the original time zone displacement.

If <time zone interval> is not specified, then no assumption is made about time zone displacement. However, should a time zone displacement be required during subsequent processing, the current default time zone displacement of the SQL-session will be applied at that time.

- 11) The truth value of a <boolean literal> is True if TRUE is specified, is False if FALSE is specified, and is Unknown if UNKNOWN is specified.

NOTE 96 — The null value of the boolean data type is equivalent to the Unknown truth value (see Subclause 4.5, “Boolean types”).

Conformance Rules

- 1) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <boolean literal>.
- 2) Without Feature F555, “Enhanced seconds precision”, in conforming SQL language, an <unsigned integer> that is a <seconds fraction> that is contained in a <timestamp literal> shall not contain more than 6 <digit>s.
- 3) Without Feature F555, “Enhanced seconds precision”, in conforming SQL language, a <time literal> shall not contain a <seconds fraction>.
- 4) Without Feature F421, “National character”, conforming SQL language shall not contain a <national character string literal>.
- 5) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval literal>.
- 6) Without Feature F271, “Compound character literals”, in conforming SQL language, a <character string literal> shall contain exactly one repetition of <character representation> (that is, it shall contain exactly one sequence of “<quote> <character representation>... <quote>”).
- 7) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <time zone interval>.
- 8) Without Feature T041, “Basic LOB data type support”, or Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary string literal>.
- 9) Without Feature T023, “Compound binary literals”, in conforming SQL language, a <binary string literal> shall contain exactly one repetition of “<quote> [{ <hexit> <hexit> }...] <quote>”.
- 10) Without Feature T024, “Spaces in binary literals”, in conforming SQL language, a <binary string literal> shall not contain a <space>.
- 11) Without Feature F393, “Unicode escapes in literals”, conforming SQL language shall not contain a <Unicode character string literal>.

5.4 Names and identifiers

*This Subclause is modified by Subclause 5.2, “Names and identifiers”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 5.2, “Names and identifiers”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 5.2, “Names and identifiers”, in ISO/IEC 9075-13.
This Subclause is modified by Subclause 5.2, “Names and identifiers”, in ISO/IEC 9075-14.*

Function

Specify names.

Format

```
<identifier> ::=
    <actual identifier>

<actual identifier> ::=
    <regular identifier>
  | <delimited identifier>
  | <Unicode delimited identifier>

<SQL language identifier> ::=
    <SQL language identifier start> [ <SQL language identifier part>... ]

<SQL language identifier start> ::=
    <simple Latin letter>

<SQL language identifier part> ::=
    <simple Latin letter>
  | <digit>
  | <underscore>

<authorization identifier> ::=
    <role name>
  | <user identifier>

<table name> ::=
    <local or schema qualified name>

<domain name> ::=
    <schema qualified name>

<schema name> ::=
    [ <catalog name> <period> ] <unqualified schema name>

<unqualified schema name> ::=
    <identifier>

<catalog name> ::=
    <identifier>

<schema qualified name> ::=
    [ <schema name> <period> ] <qualified identifier>

<local or schema qualified name> ::=
    [ <local or schema qualifier> <period> ] <qualified identifier>
```

ISO/IEC 9075-2:2008(E)
5.4 Names and identifiers

```
<local or schema qualifier> ::=  
    <schema name>  
    | <local qualifier>  
  
<qualified identifier> ::=  
    <identifier>  
  
<column name> ::=  
    <identifier>  
  
<correlation name> ::=  
    <identifier>  
  
<query name> ::=  
    <identifier>  
  
<SQL-client module name> ::=  
    <identifier>  
  
<procedure name> ::=  
    <identifier>  
  
<schema qualified routine name> ::=  
    <schema qualified name>  
  
<method name> ::=  
    <identifier>  
  
<specific name> ::=  
    <schema qualified name>  
  
<cursor name> ::=  
    <local qualified name>  
  
<local qualified name> ::=  
    [ <local qualifier> <period> ] <qualified identifier>  
  
<local qualifier> ::=  
    MODULE  
  
<host parameter name> ::=  
    <colon> <identifier>  
  
<SQL parameter name> ::=  
    <identifier>  
  
<constraint name> ::=  
    <schema qualified name>  
  
<external routine name> ::=  
    <identifier>  
    | <character string literal>  
  
<trigger name> ::=  
    <schema qualified name>  
  
<collation name> ::=  
    <schema qualified name>  
  
<character set name> ::=
```

```
[ <schema name> <period> ] <SQL language identifier>

<transliteration name> ::=
  <schema qualified name>

<transcoding name> ::=
  <schema qualified name>

<schema-resolved user-defined type name> ::=
  <user-defined type name>

<user-defined type name> ::=
  [ <schema name> <period> ] <qualified identifier>

<attribute name> ::=
  <identifier>

<field name> ::=
  <identifier>

<savepoint name> ::=
  <identifier>

<sequence generator name> ::=
  <schema qualified name>

<role name> ::=
  <identifier>

<user identifier> ::=
  <identifier>

<connection name> ::=
  <simple value specification>

<SQL-server name> ::=
  <simple value specification>

<connection user name> ::=
  <simple value specification>

<SQL statement name> ::=
  <statement name>
  | <extended statement name>

<statement name> ::=
  <identifier>

<extended statement name> ::=
  [ <scope option> ] <simple value specification>

<dynamic cursor name> ::=
  <cursor name>
  | <extended cursor name>

<extended cursor name> ::=
  [ <scope option> ] <simple value specification>

<descriptor name> ::=
  <non-extended descriptor name>
```

```
| <extended descriptor name>

<non-extended descriptor name> ::=
  <identifier>

<extended descriptor name> ::=
  [ <scope option> ] <simple value specification>

<scope option> ::=
  GLOBAL
  | LOCAL

<window name> ::=
  <identifier>
```

Syntax Rules

- 1) In an <SQL language identifier>, the number of <SQL language identifier part>s shall be less than 128.
- 2) An <SQL language identifier> is equivalent to an <SQL language identifier> in which every letter that is a lower-case letter is replaced by the corresponding upper-case letter or letters. This treatment includes determination of equivalence, representation in the Information and Definition Schemas, representation in diagnostics areas, and similar uses.

NOTE 97 — The Information Schema and Definition Schema are defined in [ISO9075-11].

- 3) An <SQL language identifier> (with every letter that is a lower-case letter replaced by the corresponding upper-case letter or letters), treated as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER, shall not be equal, according to the comparison rules in [Subclause 8.2, “<comparison predicate>”](#), to any <reserved word> (with every letter that is a lower-case letter replaced by the corresponding upper-case letter or letters), treated as the repetition of a <character string literal> that specifies a <character set specification> of SQL_IDENTIFIER.

NOTE 98 — It is the intention that no <key word> specified in ISO/IEC 9075 or revisions thereto shall end with an <underscore>.

- 4) If a <local or schema qualified name> does not contain a <local or schema qualifier>, then

Case:

- a) 04 If the <local or schema qualified name> is contained, without an intervening <schema definition>, in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then the default <unqualified schema name> for the SQL-session is implicit.
 - b) 04 If the <local or schema qualified name> is contained in a <schema definition>, then the <schema name> that is specified or implicit in the <schema definition> is implicit.
 - c) Otherwise, the <schema name> that is specified or implicit for the SQL-client module is implicit.
- 5) Let *TN* be a <table name> with a <qualified identifier> *QI* and a <local or schema qualifier> *LSQ*.

Case:

- a) 04 If *LSQ* is “MODULE”, then *TN* shall be contained in an <SQL-client module definition> *M* and the <module contents> of *M* shall contain a <temporary table declaration> *TT* whose <table name> has a <qualified identifier> equivalent to *QI*.

- b) Otherwise, *LSQ* shall be a <schema name> that identifies a schema that contains a <table definition> or <view definition> whose <table name> has a <qualified identifier> equivalent to *QI*.
- 6) If a <cursor name> *CN* with a <qualified identifier> *QI* does not contain a <local qualifier>, then the <local qualifier> *MODULE* is implicit.
- 7) 04 Let *CN* be a <cursor name>. *CN* shall be contained, without an intervening <SQL schema statement>, in an <SQL-client module definition> whose <module contents> contain a <declare cursor> or <dynamic declare cursor> whose <cursor name> is *CN*.
- 8) 04 If <user-defined type name> *UDTN* with a <qualified identifier> *QI* is specified, then

Case:

- a) If *UDTN* is simply contained in <path-resolved user-defined type name>, then

Case:

- i) If *UDTN* contains a <schema name> *SN*, then the schema identified by *SN* shall contain the descriptor of a user-defined type *UDT* such that the <qualified identifier> of *UDT* is equivalent to *QI*. *UDT* is the user-defined type identified by *UDTN*.
- ii) Otherwise,
 - 1) Case:
 - A) If *UDTN* is contained, without an intervening <schema definition>, in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then let *DP* be the SQL-path of the current SQL-session.
 - B) If *UDTN* is contained in a <schema definition>, then let *DP* be the SQL-path of that <schema definition>.
 - C) Otherwise, let *DP* be the SQL-path of the <SQL-client module definition> that contains *UDTN*.
 - 2) Let *N* be the number of <schema name>s in *DP*. Let *S_i*, 1 (one) ≤ *i* ≤ *N*, be the *i*-th <schema name> in *DP*.
 - 3) Let the *set of subject types* be the set containing every user-defined type *T* in the schema identified by some *S_i*, 1 (one) ≤ *i* ≤ *N*, such that the <qualified identifier> of *T* is equivalent to *QI*. There shall be at least one type in the set of subject types.
 - 4) Let *UDT* be the user-defined type contained in the set of subject types such that there is no other type *UDT2* for which the <schema name> of the schema that includes the user-defined type descriptor of *UDT2* precedes in *DP* the <schema name> identifying the schema that includes the user-defined type descriptor of *UDT*. *UDTN* identifies *UDT*.
 - 5) The implicit <schema name> of *UDTN* is the <schema name> of the schema that includes the user-defined type descriptor of *UDT*.
- b) If *UDTN* is simply contained in <schema-resolved user-defined type name> and *UDTN* does not contain a <schema name>, then

Case:

- i) If *UDTN* is contained, without an intervening <schema definition>, in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then the implicit <schema name> of *UDTN* is the default <unqualified schema name> for the SQL-session.
 - ii) If *UDTN* is contained in a <schema definition>, then the implicit <schema name> of *UDTN* is the <schema name> that is specified or implicit in <schema definition>.
 - iii) Otherwise, the implicit <schema name> of *UDTN* is the <schema name> that is specified or implicit in <SQL-client module definition>.
- 9) Two <user-defined type name>s are equivalent if and only if they have equivalent <qualified identifier>s and equivalent <schema name>s, regardless of whether the <schema name>s are implicit or explicit.
- 10) 11 No <unqualified schema name> shall specify DEFINITION_SCHEMA.
- 11) If a <transcoding name> does not specify a <schema name>, then INFORMATION_SCHEMA is implicit; otherwise, INFORMATION_SCHEMA shall be specified.
- 12) If a <character set name> does not specify a <schema name>, then
- Case:
- a) If <character set name> is not immediately contained in:
 - i) A <character set definition>.
 - ii) A <drop character set statement>.then <schema name> INFORMATION_SCHEMA is implicit.
 - b) Otherwise,
- Case:
- i) If the <character set name> is contained, without an intervening <schema definition>, in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then the default <unqualified schema name> for the SQL-session is implicit.
 - ii) If the <character set name> is contained in a <schema definition>, then the <schema name> that is specified or implicit in the <schema definition> is implicit.
 - iii) Otherwise, the <character set name> that is specified or implicit for the <SQL-client module definition> is implicit.
- 13) If a <schema qualified name> *SQLN* other than a <transcoding name> does not contain a <schema name>, then
- Case:
- a) If any of the following is true:
 - i) *SQLN* is immediately contained in a <collation name> that is not immediately contained in a <collation definition> or in a <drop collation statement>.
 - ii) *SQLN* is immediately contained in a <transliteration name> that is not immediately contained in a <transliteration definition> or in a <drop transliteration statement>.

then <schema name> INFORMATION_SCHEMA is implicit.

- b) If *SQL* is immediately contained in a <constraint name> that is contained in a <table definition> or an <alter table statement>, then the explicit or implicit <schema name> of the <table name> of the table identified by the <table definition> or <alter table statement> is implicit.
- c) If *SQL* is immediately contained in a <constraint name> that is contained in a <domain definition> or an <alter domain statement>, then the explicit or implicit <schema name> of the <domain name> of the domain identified by the <domain definition> or an <alter domain statement> is implicit.
- d) Otherwise,

Case:

- i) 04 If *SQL* is contained, without an intervening <schema definition>, in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then the default <unqualified schema name> for the SQL-session is implicit.
- ii) 04 If *SQL* is contained in a <schema definition>, then the <schema name> that is specified or implicit in the <schema definition> is implicit.
- iii) Otherwise, the <schema name> that is specified or implicit for the <SQL-client module definition> is implicit.

- 14) 11 If a <schema name> does not contain a <catalog name>, then

Case:

- a) If the <unqualified schema name> is contained in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then the default catalog name for the SQL-session is implicit.
- b) If the <unqualified schema name> is contained in a <module authorization clause>, then an implementation-defined <catalog name> is implicit.
- c) If the <unqualified schema name> is contained in a <schema definition> other than in a <schema name clause>, then the <catalog name> that is specified or implicit in the <schema name clause> is implicit.
- d) If the <unqualified schema name> is contained in a <schema name clause>, then

Case:

- i) If the <schema name clause> is contained in an <SQL-client module definition>, then the explicit or implicit <catalog name> contained in the <module authorization clause> is implicit.
- ii) Otherwise, an implementation-defined <catalog name> is implicit.
- e) Otherwise, the explicit or implicit <catalog name> contained in the <module authorization clause> is implicit.

- 15) Two <schema qualified name>s are equivalent if and only if their <qualified identifier>s are equivalent and their <schema name>s are equivalent, regardless of whether the <schema name>s are implicit or explicit.
- 16) Two <local or schema qualified name>s are equivalent if and only if their <qualified identifier>s are equivalent and either they both specify MODULE or they both specify or imply <schema name>s that are equivalent.

5.4 Names and identifiers

- 17) Two <character set name>s are equivalent if and only if their <SQL language identifier>s are equivalent and their <schema name>s are equivalent, regardless of whether the <schema name>s are implicit or explicit.
- 18) 13 Two <schema name>s are equivalent if and only if their <unqualified schema name>s are equivalent and their <catalog name>s are equivalent, regardless of whether the <catalog name>s are implicit or explicit.
- 19) An <identifier> that is a <correlation name> is associated with a table within a particular scope.

NOTE 99 — The scope of a <correlation name> is defined in the Syntax Rules of [Subclause 7.6](#), “<table reference>”, [Subclause 11.40](#), “<trigger definition>”, and elsewhere. Scopes may be nested. In different scopes, <correlation name>s that are equivalent <identifier>s may be associated with different tables or with the same table.

- 20) No <authorization identifier> shall specify “PUBLIC”.
- 21) Those <identifier>s that are valid <authorization identifier>s are implementation-defined.
- 22) Those <identifier>s that are valid <catalog name>s are implementation-defined.
- 23) The <data type> of <SQL-server name>, <connection name>, and <connection user name> shall be character string with an implementation-defined character set and shall have an octet length of 128 characters or less.
- 24) The <simple value specification> of <extended statement name> or <extended cursor name> shall not be a <literal>.
- 25) If an <extended statement name> or <extended cursor name> contains an <SQL parameter reference>, then it shall also contain an explicit <scope option>.
- 26) The declared type of the <simple value specification> of <extended statement name> shall be character string with an implementation-defined character set and shall have an octet length of 128 octets or less.
- 27) The declared type of the <simple value specification> of <extended cursor name> shall be character string with an implementation-defined character set and shall have an octet length of 128 octets or less.
- 28) The declared type of the <simple value specification> of <extended descriptor name> shall be character string with an implementation-defined character set and shall have an octet length of 128 octets or less.
- 29) In an <extended descriptor name>, <extended statement name>, or <extended cursor name>, if a <scope option> is not specified, then a <scope option> of LOCAL is implicit. If a <scope option> is contained in an <SQL schema statement> then it shall not contain LOCAL.
- 30) 04 If a <descriptor name> contains an <extended descriptor name> that identifies an <SQL parameter reference> and does not specify a <scope option>, then that <SQL parameter reference> is used to supply the value for the <descriptor name>.

NOTE 100 — The previous rule disambiguates between an <extended descriptor name> that is an <SQL parameter reference> and a <non-extended descriptor name> that is an <identifier> and gives precedence to the <SQL parameter reference>.

Access Rules

None.

General Rules

- 1) A <table name> identifies a table.

- 2) Within its scope, a <correlation name> identifies a table.
- 3) Within its scope, a <query name> identifies the table defined or returned by some associated <query expression body>.
- 4) A <column name> identifies a column.
- 5) A <domain name> identifies a domain.
- 6) An <authorization identifier> identifies a set of privileges.
- 7) An <SQL-client module name> identifies an SQL-client module.
- 8) A <schema qualified routine name> identifies an SQL-invoked routine.
- 9) A <method name> identifies an SQL-invoked method *M* whose descriptor is included in the schema that includes the descriptor of the user-defined type that is the type of *M*.
- 10) A <specific name> identifies an SQL-invoked routine.
- 11) A <cursor name> identifies a cursor.
- 12) A <host parameter name> identifies a host parameter.
- 13) An <SQL parameter name> identifies an SQL parameter.
- 14) An <external routine name> identifies an external routine.
- 15) A <trigger name> identifies a trigger.
- 16) A <constraint name> identifies a table constraint, a domain constraint, or an assertion.
- 17) A <catalog name> identifies a catalog.
- 18) A <schema name> identifies a schema.
- 19) A <collation name> identifies a collation.
- 20) A <character set name> identifies a character set.
- 21) A <transliteration name> identifies a character transliteration.
- 22) A <transcoding name> identifies a transcoding. All <transcoding name>s are implementation-defined.
- 23) A <connection name> identifies an SQL-connection.
- 24) A <user-defined type name> identifies a user-defined type.
- 25) An <attribute name> identifies an attribute of a structured type.
- 26) A <savepoint name> identifies a savepoint. The scope of a <savepoint name> is the SQL-transaction in which it was defined.
- 27) A <sequence generator name> identifies a sequence generator.
- 28) A <field name> identifies a field.
- 29) A <role name> identifies a role.
- 30) A <user identifier> identifies a user.

- 31) If a prepared statement *PSX* is created in SQL-session *SS* by executing a <prepare statement> *PS1* that contains an <extended statement name> *ESN1* whose value at the time of execution is *V*, then, for as long as it exists, *PSX* can be identified by an <extended statement name> *ESN2* in an <SQL procedure statement> *PS2* executed in *SS* if the value of *ESN2* at the time of execution is *V* and the <scope option> of *ESN2* is the same as the <scope option> of *ESN1*. If the <scope option> of *ESN1* is LOCAL, then *ESN2* identifies *PSX* only if *PS2* is contained in the same <SQL-client module definition> as *PS1*.

NOTE 101 — The “value at the time of execution” is defined in the General Rules of Subclause 20.6, “<prepare statement>”.

- 32) Case:

- a) A <dynamic cursor name> that is a <cursor name> is a non-extended name that identifies a declared dynamic cursor in an <SQL dynamic statement>.
- b) If a <dynamic cursor name> is an <extended cursor name>, then the value of the <simple value specification> simply contained in the <extended cursor name> is an extended name that identifies an extended dynamic cursor in an <SQL dynamic statement>.

NOTE 102 — The scope of a non-extended name is defined in Subclause 4.24.2, “Dynamic SQL statements and descriptor areas”.

- 33) A <statement name> is a non-extended name that identifies a prepared statement created by the execution of a <prepare statement>.
- 34) A <non-extended descriptor name> is a non-extended name that identifies an SQL descriptor area created by the execution of an <allocate descriptor statement>.
- 35) If an extended dynamic cursor *CSR* is created in SQL-session *SS* by executing an <allocate cursor statement> *ACS* that contains an <extended statement name> *ESN1* whose value at the time of execution is *V*, then, for as long as it exists, *CSR* can be identified by an <extended statement name> *ESN2* in an <SQL procedure statement> *PS2* executed in *SS* if the value of *ESN2* at the time of execution is *V* and the <scope option> of *ESN2* is the same as the <scope option> of *ESN1*. If the <scope option> of *ESN1* is LOCAL, then *ESN2* identifies *CSR* only if *PS2* is contained in the same <SQL-client module definition> as *ACS*.

NOTE 103 — The “value at the time of execution” is defined in the General Rules of Subclause 20.15, “<allocate cursor statement>”.

- 36) If an SQL descriptor area *SDA* is created in SQL-session *SS* by executing an <allocate descriptor statement> *ADS* that contains an <extended descriptor name> *ESN1* whose value at the time of execution is *V*, then, for as long as it exists, *SDA* can be identified by an <extended descriptor name> *ESN2* in an <SQL procedure statement> *PS2* executed in *SS* if the value of *ESN2* at the time of execution is *V* and the <scope option> of *ESN2* is the same as the <scope option> of *ESN1*. If the <scope option> of *ESN1* is LOCAL, then *ESN2* identifies *SDA* only if *PS2* is contained in the same <SQL-client module definition> as *ADS*.

NOTE 104 — The “value at the time of execution” is defined in the General Rules of Subclause 20.2, “<allocate descriptor statement>”.

- 37) A <window name> identifies a window.

Conformance Rules

- 1) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <savepoint name>.
- 2) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <role name>.
- 3) Without Feature T121, “WITH (excluding RECURSIVE) in query expression”, conforming SQL language shall not contain a <query name>.

- 4) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <attribute name>.
- 5) Without Feature T051, “Row types”, conforming SQL language shall not contain a <field name>.
- 6) Without Feature F651, “Catalog name qualifiers”, conforming SQL language shall not contain a <catalog name>.
- 7) Without Feature F771, “Connection management”, conforming SQL language shall not contain an explicit <connection name>.
- 8) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <collation name>.
- 9) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transliteration name>.
- 10) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transcoding name>.
- 11) Without Feature F821, “Local table references”, conforming SQL language shall not contain a <local or schema qualifier> that contains a <local qualifier>.
- 12) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <domain name>.
- 13) Without Feature F491, “Constraint management”, conforming SQL language shall not contain a <constraint name>.
- 14) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <character set name>.
- 15) Without Feature T601, “Local cursor references”, a <cursor name> shall not contain a <local qualifier>.
- 16) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <extended statement name> or <extended cursor name>.
- 17) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <SQL statement name>.
- 18) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain <dynamic cursor name>.
- 19) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <descriptor name>.
- 20) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window name>.
- 21) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <sequence generator name>.
- 22) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <descriptor name> that is not a <literal> or a <non-extended descriptor name>.
- 23) Without Feature B035, “Non-extended descriptor names”, conforming SQL language shall not contain a <descriptor name> that is a <non-extended descriptor name>.

(Blank page)

6 Scalar expressions

6.1 <data type>

This Subclause is modified by Subclause 6.1, “<data type>”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 6.1, “<data type>”, in ISO/IEC 9075-14.

Function

Specify a data type.

Format

```

<data type> ::=
    <predefined type>
  | <row type>
  | <path-resolved user-defined type name>
  | <reference type>
  | <collection type>

09 14 <predefined type> ::=
    <character string type> [ CHARACTER SET <character set specification> ]
      [ <collate clause> ]
  | <national character string type> [ <collate clause> ]
  | <binary string type>
  | <numeric type>
  | <boolean type>
  | <datetime type>
  | <interval type>

<character string type> ::=
    CHARACTER [ <left paren> <character length> <right paren> ]
  | CHAR [ <left paren> <character length> <right paren> ]
  | CHARACTER VARYING <left paren> <character length> <right paren>
  | CHAR VARYING <left paren> <character length> <right paren>
  | VARCHAR <left paren> <character length> <right paren>
  | <character large object type>

<character large object type> ::=
    CHARACTER LARGE OBJECT [ <left paren> <character large object length> <right paren> ]
  | CHAR LARGE OBJECT [ <left paren> <character large object length> <right paren> ]
  | CLOB [ <left paren> <character large object length> <right paren> ]

<national character string type> ::=
    NATIONAL CHARACTER [ <left paren> <character length> <right paren> ]
  | NATIONAL CHAR [ <left paren> <character length> <right paren> ]
  | NCHAR [ <left paren> <character length> <right paren> ]
  | NATIONAL CHARACTER VARYING <left paren> <character length> <right paren>

```

ISO/IEC 9075-2:2008(E)

6.1 <data type>

```
| NATIONAL CHAR VARYING <left paren> <character length> <right paren>
| NCHAR VARYING <left paren> <character length> <right paren>
| <national character large object type>

<national character large object type> ::=
    NATIONAL CHARACTER LARGE OBJECT [ <left paren> <character large object length> <right
    paren> ]
| NCHAR LARGE OBJECT [ <left paren> <character large object length> <right paren> ]
| NCLOB [ <left paren> <character large object length> <right paren> ]

<binary string type> ::=
    BINARY [ <left paren> <length> <right paren> ]
| BINARY VARYING <left paren> <length> <right paren>
| VARBINARY <left paren> <length> <right paren>
| <binary large object string type>

<binary large object string type> ::=
    BINARY LARGE OBJECT [ <left paren> <large object length> <right paren> ]
| BLOB [ <left paren> <large object length> <right paren> ]

<numeric type> ::=
    <exact numeric type>
| <approximate numeric type>

<exact numeric type> ::=
    NUMERIC [ <left paren> <precision> [ <comma> <scale> ] <right paren> ]
| DECIMAL [ <left paren> <precision> [ <comma> <scale> ] <right paren> ]
| DEC [ <left paren> <precision> [ <comma> <scale> ] <right paren> ]
| SMALLINT
| INTEGER
| INT
| BIGINT

<approximate numeric type> ::=
    FLOAT [ <left paren> <precision> <right paren> ]
| REAL
| DOUBLE PRECISION

<length> ::=
    <unsigned integer>

<character length> ::=
    <unsigned integer> [ <char length units> ]

<large object length> ::=
    <length> [ <multiplier> ]
| <large object length token>

<character large object length> ::=
    <large object length> [ <char length units> ]

<char length units> ::=
    CHARACTERS
| OCTETS

<precision> ::=
    <unsigned integer>

<scale> ::=
```



```

    <unsigned integer>

<boolean type> ::=
    BOOLEAN

<datetime type> ::=
    DATE
    | TIME [ <left paren> <time precision> <right paren> ] [ <with or without time zone> ]
    | TIMESTAMP [ <left paren> <timestamp precision> <right paren> ]
      [ <with or without time zone> ]

<with or without time zone> ::=
    WITH TIME ZONE
    | WITHOUT TIME ZONE

<time precision> ::=
    <time fractional seconds precision>

<timestamp precision> ::=
    <time fractional seconds precision>

<time fractional seconds precision> ::=
    <unsigned integer>

<interval type> ::=
    INTERVAL <interval qualifier>

<row type> ::=
    ROW <row type body>

<row type body> ::=
    <left paren> <field definition> [ { <comma> <field definition> }... ] <right paren>

<reference type> ::=
    REF <left paren> <referenced type> <right paren> [ <scope clause> ]

<scope clause> ::=
    SCOPE <table name>

<referenced type> ::=
    <path-resolved user-defined type name>

<path-resolved user-defined type name> ::=
    <user-defined type name>

<collection type> ::=
    <array type>
    | <multiset type>

<array type> ::=
    <data type> ARRAY
      [ <left bracket or trigraph> <maximum cardinality> <right bracket or trigraph> ]

<maximum cardinality> ::=
    <unsigned integer>

<multiset type> ::=
    <data type> MULTISSET

```

Syntax Rules

- 1) CHAR is equivalent to CHARACTER. DEC is equivalent to DECIMAL. INT is equivalent to INTEGER. VARCHAR is equivalent to CHARACTER VARYING. NCHAR is equivalent to NATIONAL CHARACTER. CLOB is equivalent to CHARACTER LARGE OBJECT. NCLOB is equivalent to NATIONAL CHARACTER LARGE OBJECT. VARBINARY is equivalent to BINARY VARYING. BLOB is equivalent to BINARY LARGE OBJECT.
- 2) “NATIONAL CHARACTER” is equivalent to the corresponding <character string type> with a specification of “CHARACTER SET *CSN*”, where “*CSN*” is an implementation-defined <character set name>.
- 3) If <character string type> is specified, then the collation derivation of the resulting character string type is *implicit*.

Case:

- a) If <collate clause> is specified, then the collation specified by it shall be applicable to the explicit or implicit character set *CS* of the character string type. That collation is the declared type collation of the character string type.
- b) Otherwise, the character set collation of *CS* is the declared type collation of the character string type.
- 4) The value of a <length> shall be greater than 0 (zero).
- 5) If <length> is omitted, then a <length> of 1 (one) is implicit.
- 6) If <char length units> is specified, then the character repertoire of the explicit or implicit character set of the character string type shall be UCS.
- 7) If <character length> *CL* is specified and *CL* does not contain <char length units>, then CHARACTERS is implicit.
- 8) If <large object length> is omitted, then an implementation-defined <large object length> is implicit.
- 9) The numeric value of a <large object length> *LOL* is determined as follows.

Case:

- a) If *LOL* immediately contains <unsigned integer> and does not immediately contain <multiplier>, then the numeric value of <large object length> is the numeric value of the specified <unsigned integer>.
- b) If *LOL* immediately contains <large object length token> or immediately contains <unsigned integer> and <multiplier>, then let *D* be the value of the specified <unsigned integer> or the numeric value of the sequence of <digit>s of <large object length token> interpreted as an <unsigned integer>. The numeric value of *LOL* is the numeric value resulting from the multiplication of *D* and *MS*, where *MS* is:
 - i) If <multiplier> is K, then 1,024.
 - ii) If <multiplier> is M, then 1,048,576.
 - iii) If <multiplier> is G, then 1,073,741,824.
 - iv) If <multiplier> is T, then 1,099,511,627,776.
 - v) If <multiplier> is P, then 1,125,899,906,842,624.
- 10) *LOL* shall be greater than 0 (zero).

- 11) The numeric value of a <character large object length> *CLOL* is the numeric value of the <large object length> contained in *CLOL*.
- 12) CHARACTER specifies the data type character string.
- 13) Characters in a character string are numbered beginning with 1 (one).
- 14) Case:
 - a) If the explicit or implicit <char length units> is CHARACTERS, then

Case:

 - i) If neither VARYING nor LARGE OBJECT is specified in <character string type>, then the length in characters of the character string is fixed and is the value of <length>.
 - ii) If VARYING is specified in <character string type>, then the length in characters of the character string is variable, with a minimum length of 0 (zero) and a maximum length of the value of <length>.
 - iii) If LARGE OBJECT is specified in a <character string type>, then the length in characters of the character string is variable, with a minimum length of 0 (zero) and a maximum length of the value of <large object length>.
 - b) Otherwise:
 - i) If neither VARYING nor LARGE OBJECT is specified in <character string type>, then the length in octets of the character string is fixed and is the value of <length>.
 - ii) If VARYING is specified in <character string type>, then the length in octets of the character string is variable, with a minimum length of 0 (zero) and a maximum length of the value of <length>.
 - iii) If LARGE OBJECT is specified in a <character string type>, then the length in octets of the character string is variable, with a minimum length of 0 (zero) and a maximum length of the value of <large object length>.
- 15) The maximum values of <length> and <large object length> are implementation-defined. Neither <length> nor <large object length> shall be greater than the corresponding maximum value.
- 16) If <character string type> is not contained in a <domain definition> or a <column definition> and CHARACTER SET is not specified, then an implementation-defined <character set specification> that specifies an implementation-defined character set that contains at least every character that is in <SQL language character> is implicit.

NOTE 105 — Subclause 11.25, “<domain definition>”, and Subclause 11.4, “<column definition>”, specify the result when <character string type> is contained in a <domain definition> or <column definition>, respectively.
- 17) BINARY specifies the data type binary string.
- 18) Octets in a binary string are numbered beginning with 1 (one).
- 19) Case:
 - a) If neither VARYING nor LARGE OBJECT is specified in a <binary string type>, then the length in octets of the binary string is fixed and is the value of <length>.

6.1 <data type>

- b) If VARYING is specified in a <binary string type>, then the length in octets of the binary string is variable, with a minimum of length of 0 (zero) and a maximum length of the value of <length>.
 - c) If LARGE OBJECT is specified in a <binary string type>, then the length in octets of the binary string is variable, with a minimum length of 0 (zero) and a maximum length of the value of <large object length>.
- 20) If a <precision> is omitted, then an implementation-defined <precision> is implicit.
- 21) The value of a <precision> shall be greater than 0 (zero).
- 22) If a <scale> is omitted, then a <scale> of 0 (zero) is implicit.
- 23) If an <exact numeric type> contains <scale>, then the value of <scale> shall not be greater than the value of <precision> contained in that <exact numeric type>.
- 24) For each exact numeric type *ENT*, there is an implementation-defined exact numeric type, *ENNF(ENT)*, known as the *normal form* of *ENT* (which may be *ENT* itself), such that:
- a) The name of the data type of *ENNF(ENT)* is neither DEC nor INT.
 - b) The name of the data type of *ENNF(DEC)* is the same as the name of the data type of *ENNF(DECIMAL)*.
 - c) The name of the data type of *ENNF(INT)* is the same as the name of the data type of *ENNF(INTEGER)*.
 - d) The precision, scale, and radix of *ENNF(ENT)* are the same as the precision, scale, and radix, respectively, of *ENT*.
 - e) *ENNF(ENNF(ENT))* is the same as *ENNF(ENT)*.
- 25) For the <exact numeric type>s DECIMAL and NUMERIC:
- a) The maximum value of <precision> is implementation-defined. <precision> shall not be greater than this value.
 - b) The maximum value of <scale> is implementation-defined. <scale> shall not be greater than this maximum value.
- 26) NUMERIC specifies the data type exact numeric, with the decimal precision and scale specified by the <precision> and <scale>.
- 27) 14 DECIMAL specifies the data type exact numeric, with the decimal scale specified by the <scale> and the implementation-defined decimal precision equal to or greater than the value of the specified <precision>.
- 28) SMALLINT, INTEGER, and BIGINT specify the data type exact numeric, with scale of 0 (zero) and binary or decimal precision. The choice of binary versus decimal precision is implementation-defined, but the same radix shall be chosen for all three data types. The precision of SMALLINT shall be less than or equal to the precision of INTEGER, and the precision of BIGINT shall be greater than or equal to the precision of INTEGER.
- 29) 14 FLOAT specifies the data type approximate numeric, with binary precision equal to or greater than the value of the specified <precision>. The maximum value of <precision> is implementation-defined. <precision> shall not be greater than this value.
- 30) REAL specifies the data type approximate numeric, with implementation-defined precision.
- 31) DOUBLE PRECISION specifies the data type approximate numeric, with implementation-defined precision that is greater than the implementation-defined precision of REAL.

- 32) For the <approximate numeric type>s FLOAT, REAL, and DOUBLE PRECISION, the maximum and minimum values of the exponent are implementation-defined.
- 33) For each approximate numeric type *ANT*, there is an implementation-defined approximate numeric type *ANNF(ANT)*, known as the *normal form* of *ANT* (which may be *ANT* itself), such that:
- a) The precision of *ANNF(ANT)* is the same as the precision of *ANT*.
 - b) *ANNF(ANNF(ANT))* is the same as *ANNF(ANT)*.
- 34) If <time precision> is not specified, then 0 (zero) is implicit. If <timestamp precision> is not specified, then 6 is implicit.
- 35) If <with or without time zone> is not specified, then WITHOUT TIME ZONE is implicit.
- 36) The maximum value of <time precision> and the maximum value of <timestamp precision> shall be the same implementation-defined value that is not less than 6. The values of <time precision> and <timestamp precision> shall not be greater than that maximum value.
- 37) 03 The length of a DATE is 10 positions. The length of a TIME WITHOUT TIME ZONE is 8 positions plus the <time fractional seconds precision>, plus 1 (one) position if the <time fractional seconds precision> is greater than 0 (zero). The length of a TIME WITH TIME ZONE is 14 positions plus the <time fractional seconds precision> plus 1 (one) position if the <time fractional seconds precision> is greater than 0 (zero). The length of a TIMESTAMP WITHOUT TIME ZONE is 19 positions plus the <time fractional seconds precision>, plus 1 (one) position if the <time fractional seconds precision> is greater than 0 (zero). The length of a TIMESTAMP WITH TIME ZONE is 25 positions plus the <time fractional seconds precision> plus 1 (one) position if the <time fractional seconds precision> is greater than 0 (zero).
- 38) An <interval type> specifying an <interval qualifier> whose <start field> and <end field> are both either YEAR or MONTH or whose <single datetime field> is YEAR or MONTH is a *year-month interval type*. An <interval type> that is not a year-month interval type is a *day-time interval type*.
- NOTE 106 — The length of interval data types is specified in the General Rules of Subclause 10.1, “<interval qualifier>”.
- 39) The *i*-th value of an interval data type corresponds to the *i*-th <primary datetime field>.
- 40) If <data type> is a <reference type>, then at least one of the following conditions shall be true:
- a) There exists a user-defined type descriptor whose user-defined type name is <user-defined type name> *UDTN* simply contained in <referenced type>. *UDTN* shall identify a structured type.
 - b) <reference type> is contained in the <member list> of <user-defined type definition> *UDTD* and the <path-resolved user-defined type name> simply contained in <referenced type> is equivalent to the <schema-resolved user-defined type name> contained in *UDTD*.
- 41) The <table name> contained in a <scope clause> shall identify a referenceable table whose structured type is *UDTN*.
- 42) The <table name> *STN* specified in <scope clause> identifies the scope of the reference type. This scope consists of every row in the table identified by *STN*.
- 43) An <array type> *AT* specifies an *array type*. The <data type> immediately contained in *AT* is the *element type* of the array type. The <maximum cardinality> immediately contained in *AT* is the *maximum cardinality* of a site of data type *AT*. If the maximum cardinality is not specified, then an implementation-defined maximum cardinality is implicit.

6.1 <data type>

- 44) A <multiset type> *MT* specifies a *multiset type*. The <data type> immediately contained in *MT* is the *element type* of the multiset type.
- 45) <row type> specifies the row data type.
- 46) BOOLEAN specifies the boolean data type.
- 47) If <data type> *DT1* is contained in a <data type> *DT2*, then the *root data type* of *DT1* is the outermost <data type> that contains *DT1*.

Access Rules

- 1) If <user-defined type name>, <reference type>, <row type>, or <collection type> *TY* is specified, and *TY* is usage-dependent on a user-defined type *UDT*, then

Case:

- a) If *TY* is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges shall include the USAGE privilege on *UDT*.
- b) Otherwise, the current privileges shall include the USAGE privilege on *UDT*.

General Rules

- 1) If the result of any specification or operation would be a character value one of whose characters is not in the character set of its declared type, then an exception condition is raised: *data exception — character not in repertoire*.
- 2) If any specification or operation attempts to cause an item of a character string type whose character set has a character repertoire of UCS to contain a code point that is a noncharacter, then an exception condition is raised: *data exception — noncharacter in UCS string*.
- 3) If <char length units> other than CHARACTERS is specified, then the conversion of the value of <length> to characters is implementation-defined.
- 4) For a <datetime type>,

Case:

- a) If DATE is specified, then the data type contains the <primary datetime field>s years, months, and days.
- b) If TIME is specified, then the data type contains the <primary datetime field>s hours, minutes, and seconds.
- c) If TIMESTAMP is specified, then the data type contains the <primary datetime field>s years, months, days, hours, minutes, and seconds.
- d) If WITH TIME ZONE is specified, then the data type contains the time zone datetime fields.

NOTE 107 — Within the non-null values of a <datetime type>, the value of the time zone interval is in the range −14:00 to +14:00. The range for time zone intervals is larger than many readers might expect because it is governed by political decisions in governmental bodies rather than by any natural law.

NOTE 108 — A <datetime type> contains no other fields than those specified by the preceding Rule.

- 5) For a <datetime type>, a <time fractional seconds precision> that is an explicit or implicit <time precision> or <timestamp precision> defines the number of decimal digits following the decimal point in the SECOND <primary datetime field>.
- 6) Table 9, “Valid values for datetime fields”, specifies the constraints on the values of the <primary datetime field>s in datetime values. The values of TIMEZONE_HOUR and TIMEZONE_MINUTE shall either both be non-negative or both be non-positive.

Table 9 — Valid values for datetime fields

Keyword	Valid values of datetime fields
YEAR	0001 to 9999
MONTH	01 to 12
DAY	Within the range 1 (one) to 31, but further constrained by the value of MONTH and YEAR fields, according to the rules for well-formed dates in the Gregorian calendar.
HOUR	00 to 23
MINUTE	00 to 59
SECOND	00 to 61.9(<i>N</i>) where “9(<i>N</i>)” indicates a sequence of <i>N</i> instances of the digit “9” and “ <i>N</i> ” indicates the number of digits specified by <time fractional seconds precision>.
TIMEZONE_HOUR	-14 to 14
TIMEZONE_MINUTE	-59 to 59

NOTE 109 — When the value of TIMEZONE_HOUR is either -14 or 14, the value of TIMEZONE_MINUTE is restricted to be 00 (zeros).

NOTE 110 — Datetime data types will allow dates in the Gregorian format to be stored in the date range 0001-01-01 CE through 9999-12-31 CE. The range for SECOND allows for as many as two “leap seconds”. Interval arithmetic that involves leap seconds or discontinuities in calendars will produce implementation-defined results.

- 7) An interval value can be zero, positive, or negative.
- 8) The values of the <primary datetime field>s within an interval data type are constrained as follows:
 - a) The value corresponding to the first <primary datetime field> is an integer with at most *N* digits, where *N* is the <interval leading field precision>.
 - b) Table 10, “Valid absolute values for interval fields”, specifies the constraints for the absolute values of other <primary datetime field>s in interval values.
 - c) If an interval value is zero, then all fields of the interval are zero.
 - d) If an interval value is positive, then all fields of the interval are non-negative and at least one field is positive.

6.1 <data type>

- e) If an interval value is negative, then all fields of the interval are non-positive, and at least one field is negative.

Table 10 — Valid absolute values for interval fields

Keyword	Valid values of INTERVAL fields
MONTH	0 to 11
HOURL	0 to 23
MINUTE	0 to 59
SECOND	0 to 59.9(<i>N</i>) where “9(<i>N</i>)” indicates the number of digits specified by <interval fractional seconds precision> in the <interval qualifier>.

- 9) If <data type> specifies a character string type, then a character string type descriptor is created, including the following:
 - a) The name of the data type (either CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT).
 - b) The length or maximum length in characters of the character string type.
 - c) The catalog name, schema name, and character set name of the character set of the character string type.
 - d) The catalog name, schema name, and collation name of the collation of the character string type.
- 10) If <data type> is a binary string type, then a binary string type descriptor is created, including the following:
 - a) The name of the data type (BINARY, BINARY VARYING, or BINARY LARGE OBJECT).
 - b) The length or maximum length in octets of the binary string type.
- 11) If <data type> *DT* specifies an <exact numeric type>, then a numeric data type descriptor is created for *DT*, including the following:
 - a) The name of the data type of the normal form of *DT* (NUMERIC, DECIMAL, BIGINT, INTEGER, or SMALLINT).
 - b) The precision of the normal form of *DT*.
 - c) The scale of the normal form of *DT*.
 - d) An indication of whether the precision and scale of the normal form of *DT* are expressed in decimal or binary terms.
 - e) The name of the data type specified by *DT*. If an abbreviated name was specified (INT or DEC), then the name included in the descriptor is that of corresponding full form (INTEGER or DECIMAL, respectively).
 - f) The explicit precision of *DT*, if specified.
 - g) The explicit scale of *DT*, if specified.

- 12) If <data type> *DT* specifies an <approximate numeric type>, then a numeric data type descriptor is created for *DT* including the following:
 - a) The name of the data type of the normal form of *DT* (FLOAT, REAL, or DOUBLE PRECISION).
 - b) The precision of the normal form of *DT*.
 - c) An indication that the precision is expressed in binary terms.
 - d) The name of the data type specified by *DT*.
 - e) The explicit precision of *DT*, if specified.
- 13) If <data type> specifies <boolean type>, then a boolean data type descriptor is created, including the name of the boolean type (BOOLEAN).
- 14) If <data type> specifies a <datetime type>, then a datetime data type descriptor is created, including the following:
 - a) The name of the datetime type (DATE, TIME WITHOUT TIME ZONE, TIME WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, or TIMESTAMP WITH TIME ZONE).
 - b) The value of the <time fractional seconds precision>, if DATE is not specified.
- 15) If <data type> specifies an <interval type>, then an interval data type descriptor is created, including the following:
 - a) The name of the interval data type (INTERVAL).
 - b) An indication of whether the interval data type is a year-month interval or a day-time interval.
 - c) The <interval qualifier> simply contained in the <interval type>.
- 16) If <data type> is a <collection type>, then a collection type descriptor is created. Let *KC* be the kind of collection (either ARRAY or MULTISET) specified by <collection type>. Let *ET* be the element type of <collection type>. Let *ETD* be the type designator of *ET*. The collection type descriptor includes the type designator *EDT KC*, an indication of *KC*, the descriptor of *ET*, and (in the case of array types) the maximum cardinality.
- 17) For a <row type> *RT*, the degree of *RT* is initially set to 0 (zero). The General Rules of [Subclause 6.2](#), “<field definition>”, specify the degree of *RT* during the definition of the fields of *RT*.
- 18) If the <data type> is a <row type>, then a row type descriptor is created. The row type descriptor includes a field descriptor for every <field definition> of the <row type>.
- 19) A <reference type> identifies a reference type.
- 20) 14 If <data type> is a <reference type>, then a reference type descriptor is created. Let *RDTN* be the name of the <referenced type>. The reference type descriptor includes the type designator REF(*RDTN*). If a <scope clause> is specified, then the reference type descriptor includes *STN*, identifying the scope of the reference type.

NOTE 111 — The user-defined type descriptor for a user-defined type is created in the General Rules of [Subclause 11.42](#), “<user-defined type definition>”.

Conformance Rules

- 1) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <path-resolved user-defined type name> that identifies a structured type.
- 2) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <boolean type>.
- 3) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <time precision> that does not specify 0 (zero).
- 4) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <timestamp precision> that does not specify either 0 (zero) or 6.
- 5) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval type>.
- 6) Without Feature F421, “National character”, conforming SQL language shall not contain a <national character string type>
- 7) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain <with or without time zone>.
- 8) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <reference type>.
- 9) Without Feature T051, “Row types”, conforming SQL language shall not contain a <row type>.
- 10) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array type>.
- 11) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset type>.
- 12) Without Feature S281, “Nested collection types”, conforming SQL language shall not contain a collection type that is based on a <data type> that contains a <collection type>.
- 13) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <scope clause> that is not simply contained in a <data type> that is simply contained in a <column definition>.
- 14) Without Feature S092, “Arrays of user-defined types”, conforming SQL language shall not contain an <array type> that is based on a <data type> that contains a <path-resolved user-defined type name>.
- 15) Without Feature S272, “Multisets of user-defined types”, conforming SQL language shall not contain a <multiset type> that is based on a <data type> that contains a <path-resolved user-defined type name>.
- 16) Without Feature S094, “Arrays of reference types”, conforming SQL language shall not contain an <array type> that is based on a <data type> that contains a <reference type>.
- 17) Without Feature S274, “Multisets of reference types”, conforming SQL language shall not contain a <multiset type> that is based on a <data type> that contains a <reference type>.
- 18) Without Feature S096, “Optional array bounds”, conforming SQL language shall not contain an <array type> that does not immediately contain <maximum cardinality>.
- 19) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <binary large object string type>, a <character large object type>, or a <national character large object type>.

- 20) Without Feature T061, “UCS support”, conforming SQL language shall not contain a <char length units>.
- 21) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain BIGINT.
- 22) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary string type> that is not a <binary large object string type>.

6.2 <field definition>

This Subclause is modified by Subclause 6.2, “<field definition>”, in ISO/IEC 9075-14.

Function

Define a field of a row type.

Format

```
<field definition> ::=  
  <field name> <data type>
```

Syntax Rules

- 1) Let *RT* be the <row type> that simply contains a <field definition>.
- 2) The <field name> shall not be equivalent to the <field name> of any other <field definition> simply contained in *RT*.
- 3) The declared type of the field is <data type>.
- 4) Let *DT* be the <data type>.
- 5) If *DT* is CHARACTER or CHARACTER VARYING and does not specify a <character set specification>, then the <character set specification> specified or implicit in the <schema character set specification> is implicit.

Access Rules

None.

General Rules

- 1) A data type descriptor is created that describes the declared type of the field being defined.
- 2) The degree of the row type *RT* being defined in the simply containing <row type> is increased by 1 (one).
- 3) A field descriptor is created that describes the field being defined. The field descriptor includes the following:
 - a) The <field name>.
 - b) The data type descriptor of the declared type of the field.
 - c) The ordinal position of the field in *RT*.
- 4) The field descriptor is included in the row type descriptor for *RT*.

Conformance Rules

- 1) Without Feature T051, “Row types”, conforming SQL language shall not contain a <field definition>.

6.3 <value expression primary>

This Subclause is modified by Subclause 6.3, “<value expression primary>”, in ISO/IEC 9075-14.

Function

Specify a value that is syntactically self-delimited.

Format

```

<value expression primary> ::=
    <parenthesized value expression>
  | <nonparenthesized value expression primary>

<parenthesized value expression> ::=
    <left paren> <value expression> <right paren>

14 <nonparenthesized value expression primary> ::=
    <unsigned value specification>
  | <column reference>
  | <set function specification>
  | <>window function>
  | <scalar subquery>
  | <case expression>
  | <cast specification>
  | <field reference>
  | <subtype treatment>
  | <method invocation>
  | <static method invocation>
  | <new specification>
  | <attribute or method reference>
  | <reference resolution>
  | <collection value constructor>
  | <array element reference>
  | <multiset element reference>
  | <next value expression>
  | <routine invocation>

<collection value constructor> ::=
    <array value constructor>
  | <multiset value constructor>

```

Syntax Rules

- 1) 14 The declared type of a <value expression primary> is the declared type of the simply contained <value expression>, <unsigned value specification>, <column reference>, <set function specification>, <window function>, <scalar subquery>, <case expression>, <cast specification>, <field reference>, <subtype treatment>, <method invocation>, <static method invocation>, <new specification>, <attribute or method reference>, <reference resolution>, <collection value constructor>, <array element reference>, <multiset element reference>, or <next value expression>, or the effective returns type of the simply contained <routine invocation>, respectively.

- 2) Let *NVEP* be a <nonparenthesized value expression primary> of the form “*A.B C*”, where *A* satisfies the Format of <schema name>, *B* satisfies the Format of <identifier>, and *C* satisfies the Format of <SQL argument list>. If *NVEP* satisfies the Format, Syntax Rules, and Access Rules of Subclause 6.16, “<method invocation>”, then *NVEP* is treated as a <method invocation>; otherwise, *NVEP* is treated as a <routine invocation>.

NOTE 112 — The formal grammar defined in the Format and Syntax Rules of Subclause 6.16, “<method invocation>”, and of Subclause 10.4, “<routine invocation>”, does not necessarily disambiguate between a <method invocation> and the invocation of a regular function. In such cases, the preceding Syntax Rule ensures that a <nonparenthesized value expression primary> that satisfies the Format, Syntax Rules, and Access Rules of Subclause 6.16, “<method invocation>”, is treated as a <method invocation>.

- 3) The declared type of a <collection value constructor> is the declared type of the <array value constructor> or <multiset value constructor> that it immediately contains.

Access Rules

None.

General Rules

- 1) 14 The value of a <value expression primary> is the value of the simply contained <value expression>, <unsigned value specification>, <column reference>, <set function specification>, <window function>, <scalar subquery>, <case expression>, <cast specification>, <field reference>, <subtype treatment>, <method invocation>, <static method invocation>, <new specification>, <attribute or method reference>, <reference resolution>, <collection value constructor>, <array element reference>, <multiset element reference>, <next value expression>, or <routine invocation>.
- 2) The value of a <collection value constructor> is the value of the <array value constructor> or <multiset value constructor> that it immediately contains.

Conformance Rules

None.

6.4 <value specification> and <target specification>

This Subclause is modified by Subclause 6.1, “<value specification> and <target specification>”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 6.1, “<value specification> and <target specification>”, in ISO/IEC 9075-10.

Function

Specify one or more values, host parameters, SQL parameters, dynamic parameters, or host variables.

Format

```

<value specification> ::=
    <literal>
  | <general value specification>

<unsigned value specification> ::=
    <unsigned literal>
  | <general value specification>

04 <general value specification> ::=
    <host parameter specification>
  | <SQL parameter reference>
  | <dynamic parameter specification>
  | <embedded variable specification>
  | <current collation specification>
  | CURRENT_CATALOG
  | CURRENT_DEFAULT_TRANSFORM_GROUP
  | CURRENT_PATH
  | CURRENT_ROLE
  | CURRENT_SCHEMA
  | CURRENT_TRANSFORM_GROUP_FOR_TYPE <path-resolved user-defined type name>
  | CURRENT_USER
  | SESSION_USER
  | SYSTEM_USER
  | USER
  | VALUE

04 <simple value specification> ::=
    <literal>
  | <host parameter name>
  | <SQL parameter reference>
  | <embedded variable name>

04 <target specification> ::=
    <host parameter specification>
  | <SQL parameter reference>
  | <column reference>
  | <target array element specification>
  | <dynamic parameter specification>
  | <embedded variable specification>

04 <simple target specification> ::=

```


6.4 <value specification> and <target specification>

```

    <host parameter specification>
  | <SQL parameter reference>
  | <column reference>
  | <embedded variable name>

<host parameter specification> ::=
    <host parameter name> [ <indicator parameter> ]

<dynamic parameter specification> ::=
    <question mark>

<embedded variable specification> ::=
    <embedded variable name> [ <indicator variable> ]

<indicator variable> ::=
    [ INDICATOR ] <embedded variable name>

<indicator parameter> ::=
    [ INDICATOR ] <host parameter name>

<target array element specification> ::=
    <target array reference>
        <left bracket or trigraph> <simple value specification> <right bracket or trigraph>

04 <target array reference> ::=
    <SQL parameter reference>
  | <column reference>

<current collation specification> ::=
    COLLATION FOR <left paren> <string value expression> <right paren>

```

Syntax Rules

- 1) The declared type of an <indicator parameter> shall be exact numeric with scale 0 (zero).
- 2) Each <host parameter name> shall be contained in an <SQL-client module definition>.
- 3) If USER is specified, then CURRENT_USER is implicit.

NOTE 113 — In an environment where the SQL-implementation conforms to Core SQL, conforming SQL language that contains either:

- A specified or implied <comparison predicate> that compares the <value specification> USER with a <value specification> other than USER, or
- A specified or implied assignment in which the “value” (as defined in Subclause 9.2, “Store assignment”) contains the <value specification> USER

will become non-conforming in an environment where the SQL-implementation conforms to some SQL package that supports character internationalization, unless the character repertoire of the implementation-defined character set in that environment is identical to the character repertoire of SQL_IDENTIFIER.

- 4) The declared type of CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, CURRENT_SCHEMA, and CURRENT_PATH is character string. Whether the character string is fixed length or variable length, and its length if it is fixed length or maximum length if it is variable length, are implementation-defined. The character set of the character string is SQL_IDENTIFIER. The declared type collation is the character set collation of SQL_IDENTIFIER, and the collation derivation is *implicit*.

6.4 <value specification> and <target specification>

- 5) The declared type of <string value expression> simply contained in <current collation specification> shall be character string. The declared type of <current collation specification> is character string. Whether the character string is fixed length or variable length, and its length if fixed length or maximum length if variable length, are implementation-defined. The character set of the character string is SQL_IDENTIFIER. The collation is the character set collation of SQL_IDENTIFIER, and the collation derivation is *implicit*.
- 6) The <value specification> or <unsigned value specification> VALUE shall be contained in a <domain constraint>. The declared type of an instance of VALUE is the declared type of the domain to which that domain constraint belongs.
- 7) A <target specification>, <target array reference>, or <simple target specification> that is a <column reference> shall be a new transition variable column reference.

NOTE 114 — “new transition variable column reference” is defined in Subclause 6.6, “<identifier chain>”.

- 8) If <target array element specification> is specified, then:
 - a) The declared type of the <target array reference> shall be an array type.
 - b) The declared type of a <target array element specification> is the element type of the specified <target array reference>.
 - c) The declared type of <simple value specification> shall be exact numeric with scale 0 (zero).
- 9) The declared type of an <indicator variable> shall be exact numeric with a scale of 0 (zero).
- 10) Each <embedded variable name> shall be contained in an <embedded SQL statement>.
- 11) Each <dynamic parameter specification> shall be contained in a <preparable statement> that is dynamically prepared in the current SQL-session through the execution of a <prepare statement>.
- 12) The declared type of CURRENT_DEFAULT_TRANSFORM_GROUP and of CURRENT_TRANSFORM_GROUP_FOR_TYPE <path-resolved user-defined type name> is a character string. Whether the character string is fixed length or variable length, and its length if fixed length or maximum length if variable length, are implementation-defined. The character set of the character string is SQL_IDENTIFIER. The declared type collation is the character set collation of SQL_IDENTIFIER, and the collation derivation is *implicit*.

Access Rules

None.

General Rules

- 1) A <value specification> or <unsigned value specification> specifies a value that is not selected from a table.
- 2) A <host parameter specification> identifies a host parameter or a host parameter and an indicator parameter in an <SQL-client module definition>.
- 3) 04 A <target specification> specifies a target that is a host parameter, an output SQL parameter, a column of a new transition variable, an element of a target whose declared type is an array type, a parameter used in a dynamically prepared statement, or a host variable, according to whether the <target specification> is a <host parameter specification>, an <SQL parameter reference>, a <column reference>, a <target array

6.4 <value specification> and <target specification>

element specification>, a <dynamic parameter specification>, or an <embedded variable specification>, respectively.

- 4) If a <host parameter specification> contains an <indicator parameter> and the value of the indicator parameter is negative, then the value specified by the <host parameter specification> is the null value; otherwise, the value specified by a <host parameter specification> is the value of the host parameter identified by the <host parameter name>.
- 5) The value specified by a <literal> is the value represented by that <literal>.
- 6) The value specified by CURRENT_USER is

Case:

 - a) If there is a current user identifier, then the value of that current user identifier.
 - b) Otherwise, the null value.
- 7) The value specified by SESSION_USER is the value of the SQL-session user identifier.
- 8) The value specified by CURRENT_ROLE is

Case:

 - a) If there is a current role name, then the value of that current role name.
 - b) Otherwise, the null value.
- 9) The value specified by SYSTEM_USER is equal to an implementation-defined string that represents the operating system user who executed the SQL-client module that contains the externally-invoked procedure whose execution caused the SYSTEM_USER <general value specification> to be evaluated.
- 10) The value specified by CURRENT_CATALOG is the character string that represents the current default catalog name.
- 11) The value specified by CURRENT_SCHEMA is the character string that represents the current default unqualified schema name.
- 12) The value specified by CURRENT_PATH is a <schema name list> where <catalog name>s are <delimited identifier>s and the <unqualified schema name>s are <delimited identifier>s. Each <schema name> is separated from the preceding <schema name> by a <comma> with no intervening <space>s. The schemas referenced in this <schema name list> are those referenced in the SQL-path of the current SQL-session context, in the order in which they appear in that SQL-path.
- 13) The value specified by <current collation specification> is the name of the collation of the <string value expression>.
- 14) If a <simple value specification> evaluates to the null value, then an exception condition is raised: *data exception — null value not allowed*.
- 15) 04 A <simple target specification> specifies a target that is a host parameter, an output SQL parameter, a column of a new transition variable, or a host variable, according to whether the <simple target specification> is a <host parameter specification>, an <SQL parameter reference>, a <column reference>, or an <embedded variable name>, respectively.

NOTE 115 — A <simple target specification> can never be assigned the null value.

6.4 <value specification> and <target specification>

- 16) If a <target specification> or <simple target specification> is assigned a value that is a zero-length character string, then it is implementation-defined whether an exception condition is raised: *data exception — zero-length character string*.
- 17) A <dynamic parameter specification> identifies a parameter used by a dynamically prepared statement.
- 18) An <embedded variable specification> identifies a host variable or a host variable and an indicator variable.
- 19) If an <embedded variable specification> contains an <indicator variable> and the value of the indicator variable is negative, then the value specified by the <embedded variable specification> is the null value; otherwise, the value specified by a <embedded variable specification> is the value of the host variable identified by the <embedded variable name>.
- 20) The value specified by CURRENT_DEFAULT_TRANSFORM_GROUP is the character string that represents the default transform group name in the SQL-session context.
- 21) The value specified by CURRENT_TRANSFORM_GROUP_FOR_TYPE <path-resolved user-defined type name> is the character string that represents the transform group name associated with the data type specified by <path-resolved user-defined type name>.

Conformance Rules

- 1) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <general value specification> that contains CURRENT_PATH.
- 2) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <general value specification> that contains VALUE.
- 3) Without Feature F321, “User authorization”, conforming SQL language shall not contain a <general value specification> that contains CURRENT_USER, SYSTEM_USER, or SESSION_USER.

NOTE 116 — Although CURRENT_USER and USER are semantically the same, without Feature F321, “User authorization”, CURRENT_USER shall be specified as USER.
- 4) Without Feature T332, “Extended roles”, conforming SQL language shall not contain CURRENT_ROLE.
- 5) Without Feature F762, “CURRENT_CATALOG”, conforming SQL language shall not contain a <general value specification> that contains CURRENT_CATALOG.
- 6) Without Feature F763, “CURRENT_SCHEMA”, conforming SQL language shall not contain a <general value specification> that contains CURRENT_SCHEMA.
- 7) Without Feature F611, “Indicator data types”, in conforming SQL language, the declared types of <indicator parameter>s and <indicator variable>s shall be the same implementation-defined data type.
- 8) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <general value specification> that contains a <dynamic parameter specification>.
- 9) Without Feature S097, “Array element assignment”, conforming SQL language shall not contain a <target array element specification>.
- 10) Without Feature S241, “Transform functions”, conforming SQL language shall not contain CURRENT_DEFAULT_TRANSFORM_GROUP.
- 11) Without Feature S241, “Transform functions”, conforming SQL language shall not contain CURRENT_TRANSFORM_GROUP_FOR_TYPE.

6.4 <value specification> and <target specification>

- 12) Without Feature F693, “SQL-session and client module collations”, conforming SQL language shall not contain <current collation specification>.

6.5 <contextually typed value specification>

Function

Specify a value whose data type is to be inferred from its context.

Format

```

<contextually typed value specification> ::=
    <implicitly typed value specification>
  | <default specification>

<implicitly typed value specification> ::=
    <null specification>
  | <empty specification>

<null specification> ::=
    NULL

<empty specification> ::=
    ARRAY <left bracket or trigraph> <right bracket or trigraph>
  | MULTISET <left bracket or trigraph> <right bracket or trigraph>

<default specification> ::=
    DEFAULT

```

Syntax Rules

- 1) Where the element type *ET* is determined by the context in which *ES* appears, the declared type *DT* of an <empty specification> *ES* is

Case:

- a) If *ES* simply contains ARRAY, then *ET* ARRAY[0].
- b) If *ES* simply contains MULTISET, then *ET* MULTISET.

ES is effectively replaced by CAST (*ES* AS *DT*).

NOTE 117 — In every such context, *ES* is uniquely associated with some expression or site of declared type *DT*, which thereby becomes the declared type of *ES*.

- 2) The declared type *DT* of a <null specification> *NS* is determined by the context in which *NS* appears. *NS* is effectively replaced by CAST (*NS* AS *DT*).

NOTE 118 — In every such context, *NS* is uniquely associated with some expression or site of declared type *DT*, which thereby becomes the declared type of *NS*.

- 3) The declared type *DT* of a <default specification> *DS* is the declared type of a <default option> *DO* included in some site descriptor, determined by the context in which *DS* appears. *DS* is effectively replaced by CAST (*DO* AS *DT*).

NOTE 119 — In every such context, *DS* is uniquely associated with some site of declared type *DT*, which thereby becomes the declared type of *DS*.

Access Rules

None.

General Rules

- 1) An <empty specification> specifies a collection whose cardinality is zero.
- 2) A <null specification> specifies the null value.
- 3) A <default specification> specifies the default value of some associated item.

Conformance Rules

- 1) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <empty specification> that simply contains ARRAY.
- 2) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain an <empty specification> that simply contains MULTiset.

6.6 <identifier chain>

This Subclause is modified by *Subclause 6.2*, “<identifier chain>”, in ISO/IEC 9075-4.

Function

Disambiguate a <period>-separated chain of identifiers.

Format

```
<identifier chain> ::=
  <identifier> [ { <period> <identifier> }... ]
```

```
<basic identifier chain> ::=
  <identifier chain>
```

Syntax Rules

- 1) Let IC be an <identifier chain>.
- 2) Let N be the number of <identifier>s immediately contained in IC .
- 3) Let I_i , $1 \text{ (one)} \leq i \leq N$, be the <identifier>s immediately contained in IC , in order from left to right.
- 4) Let $PIC_1 = I_1$. For each j between 2 and N , let $PIC_j = PIC_{j-1} \text{ <period> } I_j$. PIC_j is called the j -th *partial identifier chain* of IC .
- 5) Let M be the minimum of N and 4.
- 6) A column C is said to be *refinable* if the declared type of C is a row type or a structured type.
- 7) An SQL parameter P is said to be *refinable* if the declared type of P is a row type or a structured type.
- 8) 04 For at most one j between 1 (one) and M , PIC_j is called the *basis* of IC , and j is called the *basis length* of IC . The *referent* of the basis is a column C of a table or an SQL parameter SP . The basis, basis length, basis scope, and basis referent of IC are determined as follows:
 - a) If $N = 1 \text{ (one)}$, then

Case:

 - i) If IC is contained in an <order by clause> simply contained in a <query expression> QE , and the <select list> simply contained in QE directly contains a <derived column> DC whose explicit or implicit <column name> is equivalent to IC , then PIC_1 is a candidate basis, the scope of PIC_1 is QE , and the referent of PIC_1 is the column referenced by DC .
 - ii) 04 Otherwise, IC shall be contained in the scope of one or more range variables whose associated tables include a column whose <column name> is equivalent to I_1 or in the scope of a <routine name> whose associated <SQL parameter declaration list> includes an SQL parameter whose <SQL parameter name> is equivalent to I_1 . Let IS be the innermost such scope. Let the phrase *possible scope tags* denote those range variables and <routine name>s whose scope is IS .

NOTE 120 — “range variable” is defined in Subclause 4.14.6, “Operations involving tables”. It is possible for two distinct range variables to be equivalent.

Case:

- 1) If the number of possible scope is 1 (one), then let *IPST* be that possible scope tag.

Case:

- A) If *IPST* is a range variable *RV*, then let *T* be the table associated with *RV*. For every column *C* of *T* whose <column name> is equivalent to *I*₁, *PIC*₁ is a candidate basis of *IC*, the scope of *PIC*₁ is the scope of *RV*, and the referent of *PIC*₁ is *C*.

NOTE 121 — Two or more columns with equivalent column names are distinguished by their ordinal positions within *T*.

- B) 04 If *IPST* is a <routine name>, then let *SP* be the SQL parameter whose <SQL parameter name> is equivalent to *I*₁. *PIC*₁ is the basis of *IC*, the basis length is 1 (one), the basis scope is the scope of *SP*, and the basis referent is *SP*.

- 2) Otherwise, each possible scope tag shall be a range variable *RV* of a <table factor> that is directly contained in a <joined table> *JT*. *I*₁ shall be a common column name in *JT*. Let *C* be the column of *JT* that is identified by *I*₁. *PIC*₁ is a candidate basis of *IC*, the scope of *PIC*₁ is the scope of *RV*, and the referent of *PIC*₁ is *C*.

NOTE 122 — “Common column name” is defined in Subclause 7.7, “<joined table>”.

- b) If *N* > 1 (one), then the basis, basis length, basis scope, and basis referent are defined in terms of a candidate basis as follows:
 - i) If *IC* is contained in the scope of a <routine name> whose associated <SQL parameter declaration list> includes an SQL parameter *SP* whose <SQL parameter name> is equivalent to *I*₁, then *PIC*₁ is a candidate basis of *IC*, the scope of *PIC*₁ is the scope of *SP*, and the referent of *PIC*₁ is *SP*.
 - ii) If *N* = 2 and *PIC*₁ is equivalent to the <qualified identifier> of a <routine name> *RN* whose scope contains *IC* and whose associated <SQL parameter declaration list> includes an SQL parameter *SP* whose <SQL parameter name> is equivalent to *I*₂, then *PIC*₂ is a candidate basis of *IC*, the scope of *PIC*₂ is the scope of *SP*, and the referent of *PIC*₂ is *SP*.
 - iii) If *N* > 2 and *PIC*₁ is equivalent to the <qualified identifier> of a <routine name> *RN* whose scope contains *IC* and whose associated <SQL parameter declaration list> includes a refinable SQL parameter *SP* whose <SQL parameter name> is equivalent to *I*₂, then *PIC*₂ is a candidate basis of *IC*, the scope of *PIC*₂ is the scope of *SP*, and the referent of *PIC*₂ is *SP*.
 - iv) 04 If *N* = 2 and *PIC*₁ is equivalent to an exposed <correlation name> that is in scope, then let *EN* be the exposed <correlation name> that is equivalent to *PIC*₁ and has innermost scope. For every column *C* in the table associated with *EN* whose <column name> is equivalent to *I*₂ and is not a common column name, *PIC*₂ is a candidate basis of *IC*, the scope of *PIC*₂ is the scope of *EN*, and the referent of *PIC*₂ is *C*.
 - v) If *N* > 2 and *PIC*₁ is equivalent to an exposed <correlation name> that is in scope, then let *EN* be the exposed <correlation name> that is equivalent to *PIC*₁ and has innermost scope. For every

6.6 <identifier chain>

refinable column C in the table associated with EN whose <column name> is equivalent to I_2 and is not a common column name, PIC_2 is a candidate basis of IC , the scope of PIC_2 is the scope of EN , and the referent of PIC_2 is C .

- vi) If $N = 2, 3$, or 4 , and if PIC_{N-1} is equivalent to an exposed <table or query name> that is in scope, then let EN be the exposed <table or query name> that is equivalent to PIC_{N-1} and has the innermost scope. For every column C in the table associated with EN whose <column name> is equivalent to I_N and is not a common column name, PIC_N is a candidate basis of IC , the scope of PIC_N is the scope of EN , and the referent of PIC_N is C .

- c) There shall be exactly one candidate basis CB with innermost scope. The basis of IC is CB . The basis length is the length of CB . The basis scope is the scope of CB . The referent of IC is the referent of CB .

9) Let BL be the basis length of IC .

- 10) 04 If $BL < N$, then let TIC be the <value expression primary>:

```
( PICBL ) <period>
IBL+1 <period> ...
<period> IN
```

The Syntax Rules of Subclause 6.25, “<value expression>”, are applied to TIC , yielding a column reference or an SQL parameter reference, and $(N - BL)$ <field reference>s or <method invocation>s.

NOTE 123 — In this transformation, (PIC_{BL}) is interpreted as a <value expression primary> of the form <left paren> <value expression> <right paren>. PIC_{BL} is a <value expression> that is a <value expression primary> that is an <unsigned value specification> that is either a <column reference> or an <SQL parameter reference>. The identifiers I_{BL+1}, \dots, I_N are parsed using the Syntax Rules of <field reference> and <method invocation>.

- 11) A <basic identifier chain> shall be an <identifier chain> whose basis is the entire identifier chain.
- 12) A <basic identifier chain> whose basis referent is a column is a *column reference*. If the basis length is 2, and the basis scope is a <trigger definition> whose <trigger action time> is BEFORE, and I_1 is equivalent to the <new transition variable name> of the <trigger definition>, then the column reference is a *new transition variable column reference*.
- 13) 04 A <basic identifier chain> whose basis referent is an SQL parameter is an *SQL parameter reference*.
- 14) The data type of a <basic identifier chain> BIC is the data type of the basis referent of BIC .
- 15) If the declared type of a <basic identifier chain> BIC is character string, then the collation derivation of the declared type of BIC is

Case:

- a) If the declared type has a declared type collation DTC , then *implicit*.
- b) Otherwise, *none*.

Access Rules

None.

General Rules

- 1) Let *BIC* be a <basic identifier chain>.
- 2) If *BIC* is a column reference, then *BIC* references the column *C* that is the basis referent of *BIC*.
- 3) If *BIC* is an SQL parameter reference, then *BIC* references the SQL parameter *SP* of a given invocation of the SQL-invoked routine that contains *SP*.

Conformance Rules

- 1) Without Feature T325, “Qualified SQL parameter references”, conforming SQL language shall not contain an SQL parameter reference whose first <identifier> is the <qualified identifier> of a <routine name>.

6.7 <column reference>

Function

Reference a column.

Format

```
<column reference> ::=
    <basic identifier chain>
    | MODULE <period> <qualified identifier> <period> <column name>
```

Syntax Rules

- 1) Every <column reference> has a qualifying table and a qualifying scope, as defined in succeeding Syntax Rules.
- 2) A <column reference> that is a <basic identifier chain> *BIC* shall be a column reference. The qualifying scope is the basis scope of *BIC* and the qualifying table is the table that contains the basis referent of *BIC*.
- 3) If MODULE is specified, then <qualified identifier> shall be contained in an <SQL-client module definition> *M*, and shall identify a declared local temporary table *DLTT* whose <temporary table declaration> is contained in *M*, and “MODULE <period> <qualified identifier>” shall be an exposed <table or query name> *MPQI*, and <column name> shall identify a column of *DLTT*. The qualifying table is the table identified by *MPQI*, and the qualifying scope is the scope of *MPQI*.
- 4) If a <column reference> *CR* is contained in a <table expression> *TE* and the qualifying scope of *CR* contains *TE*, then *CR* is an *outer reference* to the qualifying table of *CR*.
- 5) Let *C* be the column that is referenced by *CR*. The declared type of *CR* is

Case:

 - a) If the column descriptor of *C* includes a data type, then that data type.
 - b) Otherwise, the data type identified in the domain descriptor that describes the domain that is identified by the <domain name> that is included in the column descriptor of *C*.
- 6) A <column reference> contained in a <query specification> is a *queried column reference*.
- 7) If *QCR* is a queried column reference, then:
 - a) The *qualifying query* of *QCR* is the <query specification> that simply contains the <from clause> that simply contains the <table reference> that defines the qualifying table of *QCR*.
 - b) Let *QQ* be the qualifying query of *QCR*.

Case:

 - i) If *QQ* is not grouped, or if *QCR* is contained in the <from clause> or the <where clause> simply contained in *QQ*, then *QCR* is an *ordinary column reference*.

- ii) If *QCR* is contained in the <having clause>, <window clause>, or <select list> simply contained in *QQ*, and *QCR* is contained in an aggregated argument of a <set function specification> *SFS*, and *QQ* is the aggregation query of *SFS*, then *QCR* is a *within-group-varying column reference*.
 - iii) Otherwise, *QCR* is a *group-invariant column reference*.
- 8) If *QCR* is a group-invariant column reference, then *QCR* shall be functionally dependent on the grouping columns of the qualifying query of *QCR*.

Access Rules

- 1) Let *CR* be the <column reference>.
- 2) If the qualifying table of *CR* is a base table or a viewed table, then

Case:

- a) If *CR* is contained in a <search condition> immediately contained in an <assertion definition> or a <check constraint definition>, then the applicable privileges for the <authorization identifier> that owns the containing schema shall include REFERENCES on the column referenced by *CR*.
- b) Otherwise,

Case:

- i) If *CR* is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include SELECT on the column referenced by *CR*.
- ii) Otherwise, the current privileges shall include SELECT on the column referenced by *CR*.

General Rules

- 1) Let *QCR* be a queried column reference. Let *QT* be the qualifying table of *QCR*, and let *C* be the column of *QT* that is referenced as the basis referent of *QCR*. The value of *QCR* is determined as follows:
 - a) If *QCR* is an ordinary column reference, then *QCR* denotes the value of *C* in a given row of *QT*.
 - b) If *QCR* is a within-group-varying column reference, then *QCR* denotes the values of *C* in the rows of a given group of the qualifying query of *QCR* used to construct the argument source of a <set function specification>.
 - c) If *QCR* is a group-invariant column reference, then *QCR* denotes a value that is not distinct from the value of *C* in every row of a given group of the qualifying query of *QCR*. If the most specific type of *QCR* is character string, datetime with time zone, or user-defined type, then the precise value is chosen in an implementation-dependent fashion.

Conformance Rules

- 1) Without Feature F821, “Local table references”, conforming SQL language shall not contain a <column reference> that simply contains MODULE.

6.8 <SQL parameter reference>

Function

Reference an SQL parameter.

Format

```
<SQL parameter reference> ::=  
  <basic identifier chain>
```

Syntax Rules

- 1) An <SQL parameter reference> shall be a <basic identifier chain> that is an SQL parameter reference.
- 2) The declared type of an <SQL parameter reference> is the declared type of the SQL parameter that it references.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

6.9 <set function specification>

Function

Specify a value derived by the application of a function to an argument.

Format

```
<set function specification> ::=
    <aggregate function>
  | <grouping operation>

<grouping operation> ::=
    GROUPING <left paren> <column reference>
      [ { <comma> <column reference> }... ] <right paren>
```

Syntax Rules

- 1) If <aggregate function> specifies a <general set function>, then the <value expression> simply contained in the <general set function> shall not contain a <set function specification> or a <query expression>.
- 2) If <aggregate function> specifies <binary set function>, then neither the <dependent variable expression> nor the <independent variable expression> simply contained in the <binary set function> shall contain a <set function specification> or a <query expression>.
- 3) A <value expression> *VE* simply contained in a <set function specification> *SFE* is an *aggregated argument* of *SFE* if either *SFE* is not an <ordered set function> or *VE* is simply contained in a <within group specification>; otherwise, *VE* is a *non-aggregated argument* of *SFE*.
- 4) A column reference *CR* contained in an aggregated argument of a <set function specification> *SFS* is called an *aggregated column reference* of *SFS*.
- 5) If <aggregate function> specifies a <filter clause>, then the <search condition> immediately contained in <filter clause> shall not contain a <set function specification>.
- 6) The *aggregation query* of a <set function specification> *SFS* is determined as follows.
Case:
 - a) If *SFS* has no aggregated column reference, then the aggregation query of *SFS* is the innermost <query specification> that contains *SFS*.
 - b) Otherwise, the innermost qualifying query of the aggregated column references of *SFS* is the aggregation query of *SFS*.
- 7) *SFS* shall be contained in the <having clause>, <window clause>, or <select list> of its aggregation query.
- 8) Let *CR* be an aggregated column reference of *SFS* such that the qualifying query *QQ* of *CR* is not the aggregation query of *SFS*. If *QQ* is grouped and *SFS* is contained in the <having clause>, <window clause>, or <select list> of *QQ*, then *CR* shall be functionally dependent on the grouping columns of *QQ*.
- 9) If <aggregate function> is specified, then the declared type of the result is the declared type of the <aggregate function>.

10) If a <grouping operation> is specified, then:

- a) Let T be the argument source of <set function specification> that contains <grouping operation>. Each <column reference> shall reference a grouping column of T .
- b) The declared type of the result is exact numeric with an implementation-defined precision and a scale of 0 (zero).
- c) If more than one <column reference> is specified, then let N be the number of <column reference>s and let CR_i , $1 \text{ (one)} \leq i \leq N$, be the i -th <column reference>.

GROUPING (CR_1 , ..., CR_{N-1} , CR_N)

is equivalent to:

(2 * GROUPING (CR_1 , ..., CR_{N-1}) + GROUPING (CR_N))

Access Rules

None.

General Rules

- 1) If <aggregate function> is specified, then the result is the value of the <aggregate function>.

NOTE 124 — The value of <grouping operation> is computed by means of syntactic transformations defined in [Subclause 7.9](#), “<group by clause>”.

Conformance Rules

- 1) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <grouping operation>.
- 2) Without Feature T433, “Multiargument GROUPING function”, conforming SQL language shall not contain a <grouping operation> that contains more than one <column reference>.

6.10 <window function>

Function

Specify a window function.

Format

```
<window function> ::=
    <window function type> OVER <window name or specification>

<window function type> ::=
    <rank function type> <left paren> <right paren>
    | ROW_NUMBER <left paren> <right paren>
    | <aggregate function>

<rank function type> ::=
    RANK
    | DENSE_RANK
    | PERCENT_RANK
    | CUME_DIST

<window name or specification> ::=
    <window name>
    | <in-line window specification>

<in-line window specification> ::=
    <window specification>
```

Syntax Rules

- 1) An <aggregate function> simply contained in a <window function> shall not simply contain a <hypothetical set function>.
- 2) Let *OF* be the <window function>.
- 3) Case:
 - a) If *OF* is contained in an <order by clause>, then the <order by clause> shall be simply contained in a <query expression> that is a simple table query. Let *TE* be the <table expression> contained in the result of the syntactic transformation of a simple table query, as specified in Subclause 7.13, “<query expression>”.
 - b) Otherwise, *OF* shall be contained in a <select list> that is immediately contained in a <query specification> *QS* or a <select statement: single row> *SSSR*. Let *QSS* be the innermost <query specification> contained in *QS* that contains *OF*. Let *TE* be the <table expression> immediately contained in *QSS* or *SSSR*.
- 4) *OF* shall not contain an outer reference or a <query expression>.
- 5) Let *WNS* be the <window name or specification>. Let *WDX* be a window structure descriptor that describes the window defined by *WNS*.
- 6) If <rank function type> or ROW_NUMBER is specified, then:

- a) If RANK or DENSE_RANK is specified, then the window ordering clause *WOC* of *WDX* shall be present.
- b) The window framing of *WDX* shall not be present.
- c) Case:
 - i) If *WNS* is a <window name>, then let *WNS1* be *WNS*.
 - ii) Otherwise, let *WNS1* be the <window specification details> contained in *WNS*.
- d) RANK () OVER *WNS* is equivalent to:

```
( COUNT ( * ) OVER ( WNS1 RANGE UNBOUNDED PRECEDING )
- COUNT ( * ) OVER ( WNS1 RANGE CURRENT ROW ) + 1 )
```

- e) If DENSE_RANK is specified, then:
 - i) Let VE_1, \dots, VE_N be an enumeration of the <value expression>s that are <sort key>s simply contained in *WOC*.
 - ii) DENSE_RANK () OVER *WNS* is equivalent to the <window function>:

```
COUNT ( DISTINCT ROW (  $VE_1, \dots, VE_N$  ) )
OVER ( WNS1 RANGE UNBOUNDED PRECEDING )
```

- f) ROW_NUMBER () OVER *WNS* is equivalent to the <window function>:

```
COUNT ( * ) OVER ( WNS1 ROWS UNBOUNDED PRECEDING )
```

- g) Let *ANT1* be an approximate numeric type with implementation-defined precision. PERCENT_RANK () OVER *WNS* is equivalent to:

```
CASE
  WHEN COUNT ( * ) OVER ( WNS1 RANGE BETWEEN UNBOUNDED PRECEDING
                        AND UNBOUNDED FOLLOWING ) = 1
  THEN CAST ( 0 AS ANT1 )
  ELSE
    ( CAST ( RANK ( ) OVER ( WNS1 ) AS ANT1 ) - 1 ) /
    ( COUNT ( * ) OVER ( WNS1 RANGE BETWEEN UNBOUNDED PRECEDING
                        AND UNBOUNDED FOLLOWING ) - 1 )
END
```

- h) Let *ANT2* be an approximate numeric type with implementation-defined precision. CUME_DIST () OVER *WNS* is equivalent to:

```
( CAST ( COUNT ( * ) OVER
  ( WNS1 RANGE UNBOUNDED PRECEDING ) AS ANT2 ) /
COUNT ( * ) OVER ( WNS1 RANGE BETWEEN UNBOUNDED PRECEDING
                  AND UNBOUNDED FOLLOWING ) )
```

- 7) Let *SL* be the <select list> that simply contains *OF*.

NOTE 125 — If *OF* is originally contained in an <order by clause> of a <query expression> that is a simple table query, then the syntactic transformation of a simple table query, as specified in Subclause 7.13, “<query expression>”, is applied prior to this rule.

- 8) Let *SQ* be the <set quantifier> of the <query specification> or <select statement: single row> that simply contains *SL*. If there is no <set quantifier>, then let *SQ* be a zero-length string.
- 9) If <in-line window specification> is specified, then:
 - a) Let *WS* be the <window specification>.
 - b) Let *WSN* be an implementation-dependent <window name> that is not equivalent to any other <window name> in the <table expression> or <select statement: single row> that simply contains *WS*.
 - c) Let *OFT* be the <window function type>.
 - d) Let *SLNEW* be the <select list> that is obtained from *SL* by replacing *OF* by:


```
OFT OVER WSN
```
 - e) Let *FC*, *WC*, *GBC*, and *HC* be <from clause>, <where clause>, <group by clause>, and <having clause>, respectively, of *TE*. If any of <where clause>, <group by clause>, or <having clause> is missing, then let *WC*, *GBC*, or *HC*, respectively, be a zero-length string.
 - f) Case:
 - i) If there is no <window clause> immediately contained in *TE*, then let *WICNEW* be:


```
WINDOW WSN AS WS
```
 - ii) Otherwise, let *WIC* be the <window clause> immediately contained in *TE* and let *WICNEW* be:


```
WIC, WSN AS WS
```
 - g) Let *TENNEW* be:


```
FC WC GBC HC WICNEW
```
 - h) Case:
 - i) If *OF* is simply contained in a <query specification>, then that <query specification> is equivalent to:


```
SELECT SQ SLNEW TENNEW
```
 - ii) Otherwise, *OF* is simply contained in a <select statement: single row>. Let *STL* be the <select target list> of that <select statement: single row>. The <select statement: single row> is equivalent to:


```
SELECT SQ SLNEW INTO STL TENNEW
```
- 10) If the window ordering clause or the window framing clause of the window structure descriptor that describes the <window name or specification> is present, then no <aggregate function> simply contained in <window function> shall specify **DISTINCT** or <ordered set function>.

Access Rules

None.

General Rules

- 1) The value of <window function> is the value of the <aggregate function>.

Conformance Rules

- 1) Without Feature T611, “Elementary OLAP operations”, conforming SQL language shall not contain a <window function>.
- 2) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window name>.
- 3) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain PERCENT_RANK or CUME_DIST.
- 4) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window function> that simply contains ROW_NUMBER and immediately contains a <window name or specification> whose window structure descriptor does not contain a window ordering clause.

6.11 <case expression>

This Subclause is modified by Subclause 6.4, “<case expression>”, in ISO/IEC 9075-14.

Function

Specify a conditional value.

Format

```

<case expression> ::=
    <case abbreviation>
    | <case specification>

<case abbreviation> ::=
    NULLIF <left paren> <value expression> <comma> <value expression> <right paren>
    | COALESCE <left paren> <value expression>
      { <comma> <value expression> }... <right paren>

<case specification> ::=
    <simple case>
    | <searched case>

<simple case> ::=
    CASE <case operand> <simple when clause>... [ <else clause> ] END

<searched case> ::=
    CASE <searched when clause>... [ <else clause> ] END

<simple when clause> ::=
    WHEN <when operand list> THEN <result>

<searched when clause> ::=
    WHEN <search condition> THEN <result>

<else clause> ::=
    ELSE <result>

<case operand> ::=
    <row value predicand>
    | <overlaps predicate part 1>

<when operand list> ::=
    <when operand> [ { <comma> <when operand> }... ]

14 <when operand> ::=
    <row value predicand>
    | <comparison predicate part 2>
    | <between predicate part 2>
    | <in predicate part 2>
    | <character like predicate part 2>
    | <octet like predicate part 2>
    | <similar predicate part 2>
    | <regex like predicate part 2>
    | <null predicate part 2>
    | <quantified comparison predicate part 2>

```

6.11 <case expression>

```

| <normalized predicate part 2>
| <match predicate part 2>
| <overlaps predicate part 2>
| <distinct predicate part 2>
| <member predicate part 2>
| <submultiset predicate part 2>
| <set predicate part 2>
| <type predicate part 2>

<result> ::=
    <result expression>
  | NULL

<result expression> ::=
    <value expression>

```

Syntax Rules

- 1) If a <case expression> specifies a <case abbreviation>, then:
 - a) A <value expression> generally contained in the <case abbreviation> shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that is possibly non-deterministic or that possibly modifies SQL-data.
 - b) NULLIF (V_1 , V_2) is equivalent to the following <case specification>:

```

CASE WHEN
   $V_1 = V_2$  THEN
  NULL ELSE  $V_1$ 
END

```

- c) COALESCE (V_1 , V_2) is equivalent to the following <case specification>:

```

CASE WHEN  $V_1$ 
  IS NOT NULL THEN
   $V_1$  ELSE
   $V_2$  END

```

- d) COALESCE (V_1 , V_2 , . . . , V_n), for $n \geq 3$, is equivalent to the following <case specification>:

```

CASE WHEN  $V_1$ 
  IS NOT NULL THEN
   $V_1$  ELSE COALESCE
  (  $V_2$ , . . . ,
   $V_n$  ) END

```

- 2) If a <case specification> specifies a <simple case>, then let *CO* be the <case operand>.
 - a) *CO* shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that is possibly non-deterministic or that possibly modifies SQL-data.
 - b) If *CO* is <overlaps predicate part 1>, then each <when operand> shall be <overlaps predicate part 2>. If *CO* is <row value predicand>, then no <when operand> shall be an <overlaps predicate part 2>.

- c) Let N be the number of <simple when clause>s.
- d) For each i between 1 (one) and N , let WOL_i be the <when operand list> of the i -th <simple when clause>. Let $M(i)$ be the number of <when operand>s simply contained in WOL_i . For each j between 1 and $M(i)$, let $WO_{i,j}$ be the j -th <when operand> simply contained in WOL_i .
- e) For each i between 1 (one) and N , and for each j between 1 (one) and $M(i)$,

Case:

- i) If $WO_{i,j}$ is a <row value predicand>, then let $EWO_{i,j}$ be

$$= WO_{i,j}$$

- ii) Otherwise, let $EWO_{i,j}$ be $WO_{i,j}$.

- f) Let R_i be the <result> of the i -th <simple when clause>.
- g) If <else clause> is specified, then let $CEEC$ be that <else clause>; otherwise, let $CEEC$ be a character string of length 0 (zero).
- h) The <simple case> is equivalent to a <searched case> in which the i -th <searched when clause> takes the form:

```
WHEN ( CO EWOi,1 ) OR
. . . OR
( CO EWOi,M(i) )
THEN Ri
```

- i) The <else clause> of the equivalent <searched case> takes the form:

$CEEC$

- 3) At least one <result> in a <case specification> shall specify a <result expression>.
- 4) If an <else clause> is not specified, then ELSE NULL is implicit.
- 5) The declared type of a <case specification> is determined by applying [Subclause 9.3, “Result of data type combinations”](#), to the declared types of all <result expression>s in the <case specification>.

Access Rules

None.

General Rules

- 1) Case:
 - a) If a <result> specifies NULL, then its value is the null value.
 - b) If a <result> specifies a <value expression>, then its value is the value of that <value expression>.
- 2) Case:

6.11 <case expression>

- a) If the value of the <search condition> of some <searched when clause> in a <case specification> is True, then the value of the <case expression> is the value of the <result> of the first (leftmost) <searched when clause> whose <search condition> evaluates to True, cast as the declared type of the <case specification>.
- b) If no <search condition> in a <case specification> evaluates to True, then the value of the <case expression> is the value of the <result> of the explicit or implicit <else clause>, cast as the declared type of the <case specification>.

Conformance Rules

- 1) Without Feature F262, “Extended CASE expression”, in conforming SQL language, a <case operand> immediately contained in a <simple case> shall be a <row value predicand> that is a <row value constructor predicand> that is a single <common value expression> or <boolean predicand>.
- 2) Without Feature F262, “Extended CASE expression”, in conforming SQL language, a <when operand> contained in a <simple when clause> shall be a <row value predicand> that is a <row value constructor predicand> that is a single <common value expression> or <boolean predicand>.
- 3) Without Feature F263, “Comma-separated predicates in simple CASE expression”, in conforming SQL language, a <when operand list> contained in a <simple when clause> shall simply contain exactly one <when operand>.

6.12 <cast specification>

This Subclause is modified by Subclause 6.2, “<cast specification>”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 6.5, “<cast specification>”, in ISO/IEC 9075-14.

Function

Specify a data conversion.

Format

```

14 <cast specification> ::=
    CAST <left paren> <cast operand> AS <cast target> <right paren>

<cast operand> ::=
    <value expression>
  | <implicitly typed value specification>

<cast target> ::=
    <domain name>
  | <data type>

```

Syntax Rules

- 1) Case:
 - a) If a <domain name> is specified, then let *TD* be the data type of the specified domain.
 - b) If a <data type> is specified, then let *TD* be the data type identified by <data type>. <data type> shall not contain a <collate clause>.
- 2) 09 The declared type of the result of the <cast specification> is *TD*.
- 3) 14 If the <cast operand> is a <value expression>, then let *SD* be the declared type of the <value expression>.
- 4) Let *C* be some column and let *CO* be the <cast operand> of a <cast specification> *CS*. *C* is a *leaf column* of *CS* if *CO* consists of a single column reference that identifies *C* or of a single <cast specification> *CS1* of which *C* is a leaf column.
- 5) If the <cast operand> specifies an <empty specification>, then *TD* shall be a collection type.
- 6) 09 14 If the <cast operand> is a <value expression>, then the valid combinations of *TD* and *SD* in a <cast specification> are given by the following table. “Y” indicates that the combination is syntactically valid without restriction; “M” indicates that the combination is valid subject to other Syntax Rules in this Subclause being satisfied; and “N” indicates that the combination is not valid:

<i>SD</i>	<i>TD</i>														
	EN	AN	C	D	T	TS	YM	DT	BO	UDT	B	RT	CT	RW	
EN	Y	Y	Y	N	N	N	M	M	N	M	N	M	N	N	
AN	Y	Y	Y	N	N	N	N	N	N	M	N	M	N	N	
C	Y	Y	Y	Y	Y	Y	Y	Y	Y	M	N	M	N	N	
D	N	N	Y	Y	N	Y	N	N	N	M	N	M	N	N	
T	N	N	Y	N	Y	Y	N	N	N	M	N	M	N	N	

6.12 <cast specification>

TS	N	N	Y	Y	Y	Y	N	N	N	M	N	M	N	N
YM	M	N	Y	N	N	N	Y	N	N	M	N	M	N	N
DT	M	N	Y	N	N	N	N	Y	N	M	N	M	N	N
BO	N	N	Y	N	N	N	N	N	Y	M	N	M	N	N
UDT	M	M	M	M	M	M	M	M	M	M	M	M	N	N
B	N	N	N	N	N	N	N	N	N	M	Y	M	N	N
RT	M	M	M	M	M	M	M	M	M	M	M	M	N	N
CT	N	N	N	N	N	N	N	N	N	N	N	N	M	N
RW	N	N	N	N	N	N	N	N	N	N	N	N	N	M

Where:

EN	=	Exact Numeric
AN	=	Approximate Numeric
C	=	Character (Fixed- or Variable-Length, or Character Large Object)
D	=	Date
T	=	Time
TS	=	Timestamp
YM	=	Year-Month Interval
DT	=	Day-Time Interval
BO	=	Boolean
UDT	=	User-Defined Type
B	=	Binary (Fixed- or Variable-Length or Binary Large Object)
RT	=	Reference type
CT	=	Collection type
RW	=	Row type

- 7) If *TD* is an interval and *SD* is exact numeric, then *TD* shall contain only a single <primary datetime field>.
- 8) If *TD* is exact numeric and *SD* is an interval, then *SD* shall contain only a single <primary datetime field>.
- 9) If *SD* is character string and *TD* is fixed-length, variable-length, or large object character string, then the character repertoires of *SD* and *TD* shall be the same.
- 10) If *TD* is a fixed-length, variable-length, or large object character string, then the declared type collation of the <cast specification> is the character set collation of the character set of *TD* and its collation derivation is *implicit*.
- 11) If the <cast operand> is a <value expression> and either *SD* or *TD* is a user-defined type, then either *TD* shall be a supertype of *SD* or there shall be a data type *P* such that all of the following are satisfied:
 - a) The type designator of *P* is in the type precedence list of *SD*.
 - b) There is a user-defined cast CF_P whose user-defined cast descriptor includes *P* as the source data type and *TD* as the target data type.
 - c) The type designator of no other data type *Q* that is included as the source data type in the user-defined cast descriptor of some user-defined cast CF_Q that has *TD* as the target data type precedes the type designator of *P* in the type precedence list of *SD*.
- 12) If the <cast operand> is a <value expression> and either *SD* or *TD* is a reference type, then:
 - a) Let *RTSD* and *RTTD* be the referenced types of *SD* and *TD*, respectively.
 - b) If <data type> is specified and contains a <scope clause>, then let *STD* be that scope. Otherwise, let *STD*, possibly empty, be the scope included in the reference type descriptor of *SD*.
 - c) Either *RSTD* and *RTTD* shall be compatible, or there shall be a data type *P* in the type precedence list of *SD* such that all of the following are satisfied:

- i) There is a user-defined cast CF_P whose user-defined cast descriptor includes P as the source data type and TD as the target data type.
- ii) The type designator of no other data type Q that is included as the source data type in the user-defined cast descriptor of some user-defined cast CF_Q that has TD as the target data type precedes the type designator of P in the type precedence list of SD .

13) If SD is a collection type, then:

- a) Let ESD be the element type of SD .
- b) Let ETD be the element type of TD .

$CAST (VALUE AS ETD)$

where $VALUE$ is a <value expression> of declared type ESD , shall be a valid <cast specification>.

14) If SD is a row type, then:

- a) Let DSD be the degree of SD .
- b) Let DTD be the degree of TD .
- c) DSD shall be equal to DTD .
- d) Let FSD_i and FTD_i , $1 \text{ (one)} \leq i \leq DSD$, be the i -th field of SD and TD , respectively.
- e) Let $TFSD_i$ and $TFTD_i$, $1 \text{ (one)} \leq i \leq DSD$, be the declared type of FSD_i and the declared type of FTD_i , respectively.
- f) For i varying from 1 (one) to DSD , the <cast specification>:

$CAST (VALUE_i AS$
 $TFTD_i)$

where $VALUE_i$ is an arbitrary <value expression> of declared type $TFSD_i$, shall be a valid <cast specification>.

15) 14 If <domain name> is specified, then let D be the domain identified by the <domain name>. The schema identified by the explicit or implicit qualifier of the <domain name> shall include the descriptor of D .

Access Rules

1) If <domain name> is specified, then

Case:

- a) If <cast specification> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include USAGE on the domain identified by <domain name>.
 - b) Otherwise, the current privileges shall include USAGE on the domain identified by <domain name>.
- 2) If the <cast operand> is a <value expression> and either SD or TD is a user-defined type or a reference type, then

Case:

- a) If <cast specification> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include EXECUTE on CF_P .
- b) Otherwise, the current privileges shall include EXECUTE on CF_P .

General Rules

- 1) Let CS be the <cast specification>. If the <cast operand> is a <value expression> VE , then let SV be the value of VE .
- 2) Case:
 - a) If the <cast operand> specifies NULL, then the result of CS is the null value and no further General Rules of this Subclause are applied.
 - b) If the <cast operand> specifies an <empty specification>, then the result of CS is an empty collection of declared type TD and no further General Rules of this Subclause are applied.
 - c) If SV is the null value, then the result of CS is the null value and no further General Rules of this Subclause are applied.

- 3) 14 If either SD or TD is a user-defined type, then

Case:

- a) If TD is a supertype of SD , then TV is SV .
- b) Otherwise:
 - i) Let CP be the cast function contained in the user-defined cast descriptor of CF_P .
 - ii) Let SAL be a static SQL argument list that has a single SQL-argument that is <value expression>. The General Rules of Subclause 10.4, "<routine invocation>", are applied with SAL as *STATIC SQL ARG LIST* and CP as *SUBJECT ROUTINE*, yielding a value TR that is the result of the invocation of CP .
 - iii) Case:
 - 1) If TD is a user-defined type, then TV is TR .
 - 2) Otherwise, TV is the result of

$CAST (TR AS TD)$

- 4) If either SD or TD is a reference type, then

Case:

- a) If $RSTD$ and $RTTD$ are compatible, then:
 - i) TV is SV .
 - ii) The scope in the reference type descriptor of TV is STD .

b) Otherwise:

- i) Let *CP* be the cast function contained in the user-defined cast descriptor of *CF_p*.
- ii) Let *SAL* be a static SQL argument list that has a single SQL-argument that is <value expression>. The General Rules of Subclause 10.4, “<routine invocation>”, are applied with *SAL* as *STATIC SQL ARG LIST* and *CP* as *SUBJECT ROUTINE*, yielding value *TV* that is the result of the invocation of *CP*.
- iii) The scope in the reference type descriptor of *TV* is *STD*.

5) If *SD* is an array type, then:

- a) Let *SC* be the cardinality of *SV*.
- b) For *i* varying from 1 (one) to *SC*, the following <cast specification> is applied:

CAST (*VE*[*i*] AS *ETD*)

yielding value *TVE_i*.

- c) If *TD* is an array type, then let *TC* be the maximum cardinality of *TD*.

Case:

- i) If *SC* is greater than *TC*, then an exception condition is raised: *data exception — array data, right truncation*.
- ii) Otherwise, *TV* is the array with elements *TVE_i*, 1 (one) ≤ *i* ≤ *SC*.

- d) If *TD* is a multiset type, then *TV* is the multiset with elements *TVE_i*, 1 (one) ≤ *i* ≤ *SC*.

6) If *SD* is a multiset type, then:

- a) Let *SC* be the cardinality of *SV*.
- b) If *TD* is an array type, then let *TC* be the maximum cardinality of *TD*.

Case:

- i) If *SC* is greater than *TC*, then an exception condition is raised: *data exception — array data, right truncation*.
- ii) Otherwise, *TV* is the array resulting from the evaluation of

ARRAY ((SELECT CAST (*M.E* AS *ETD*) FROM UNNEST (*VE*) AS *M*(*E*)))

- c) If *TD* is a multiset type, then *TV* is the multiset resulting from the evaluation of

MULTISET ((SELECT CAST (*M.E* AS *ETD*) FROM UNNEST (*VE*) AS *M*(*E*)))

7) If *SD* is a row type, then *TV* is the row resulting from the evaluation of

ROW (CAST (*VE.FSD₁* AS *TFTD₁*) ,
 CAST (*VE.FSD₂* AS *TFTD₂*) ,
 ...
 CAST (*VE.FSD_{DSD}* AS *TFTD_{DSD}*))

6.12 <cast specification>

8) If *TD* is exact numeric, then

Case:

a) If *SD* is exact numeric or approximate numeric, then

Case:

- i) If there is a representation of *SV* in the data type *TD* that does not lose any leading significant digits after rounding or truncating if necessary, then *TV* is that representation. The choice of whether to round or truncate is implementation-defined.
 - ii) Otherwise, an exception condition is raised: *data exception — numeric value out of range*.
- b) If *SD* is character string, then *SV* is replaced by *SV* with any leading or trailing <space>s removed.

Case:

- i) If *SV* does not comprise a <signed numeric literal> as defined by the rules for <literal> in Subclause 5.3, “<literal>”, then an exception condition is raised: *data exception — invalid character value for cast*.
- ii) Otherwise, let *LT* be that <signed numeric literal>. The <cast specification> is equivalent to

CAST (*LT* AS *TD*)

c) If *SD* is an interval data type, then

Case:

- i) If there is a representation of *SV* in the data type *TD* that does not lose any leading significant digits, then *TV* is that representation.
- ii) Otherwise, an exception condition is raised: *data exception — numeric value out of range*.

9) If *TD* is approximate numeric, then

Case:

a) If *SD* is exact numeric or approximate numeric, then

Case:

- i) If there is a representation of *SV* in the data type *TD* that does not lose any leading significant digits after rounding or truncating if necessary, then *TV* is that representation. The choice of whether to round or truncate is implementation-defined.
 - ii) Otherwise, an exception condition is raised: *data exception — numeric value out of range*.
- b) If *SD* is character string, then *SV* is replaced by *SV* with any leading or trailing <space>s removed.

Case:

- i) If *SV* does not comprise a <signed numeric literal> as defined by the rules for <literal> in Subclause 5.3, “<literal>”, then an exception condition is raised: *data exception — invalid character value for cast*.
- ii) Otherwise, let *LT* be that <signed numeric literal>. The <cast specification> is equivalent to

CAST (LT AS TD)

10) If *TD* is fixed-length character string, then let *LTD* be the length in characters of *TD*.

Case:

a) If *SD* is exact numeric, then:

- i) Let *YP* be the shortest character string that conforms to the definition of <exact numeric literal> in Subclause 5.3, “<literal>”, whose scale is the same as the scale of *SD* and whose interpreted value is the absolute value of *SV*.
- ii) Case:
 - 1) If *SV* is less than 0 (zero), then let *Y* be the result of ' - ' || *YP*.
 - 2) Otherwise, let *Y* be *YP*.
- iii) Case:
 - 1) If *Y* contains any <SQL language character> that is not in the character repertoire of *TD*, then an exception condition is raised: *data exception — invalid character value for cast*.
 - 2) If the length in characters *LY* of *Y* is equal to *LTD*, then *TV* is *Y*.
 - 3) If the length in characters *LY* of *Y* is less than *LTD*, then *TV* is *Y* extended on the right by *LTD*–*LY* <space>s.
 - 4) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.

b) If *SD* is approximate numeric, then:

- i) Let *YP* be a character string as follows.
Case:
 - 1) If *SV* equals 0 (zero), then *YP* is '0E0'.
 - 2) Otherwise, *YP* is the shortest character string that conforms to the definition of <approximate numeric literal> in Subclause 5.3, “<literal>”, whose interpreted value is equal to the absolute value of *SV* and whose <mantissa> consists of a single <digit> that is not '0' (zero), followed by a <period> and an <unsigned integer>.
- ii) Case:
 - 1) If *SV* is less than 0 (zero), then let *Y* be the result of ' - ' || *YP*.
 - 2) Otherwise, let *Y* be *YP*.
- iii) Case:
 - 1) If *Y* contains any <SQL language character> that is not in the character repertoire of *TD*, then an exception condition is raised: *data exception — invalid character value for cast*.
 - 2) If the length in characters *LY* of *Y* is equal to *LTD*, then *TV* is *Y*.
 - 3) If the length in characters *LY* of *Y* is less than *LTD*, then *TV* is *Y* extended on the right by *LTD*–*LY* <space>s.

4) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.

- c) If *SD* is fixed-length character string, variable-length character string, or large object character string, then

Case:

- i) If the length in characters of *SV* is equal to *LTD*, then *TV* is *SV*.
 - ii) If the length in characters of *SV* is larger than *LTD*, then *TV* is the first *LTD* characters of *SV*. If any of the remaining characters of *SV* are non-`<space>` characters, then a completion condition is raised: *warning — string data, right truncation*.
 - iii) If the length in characters *M* of *SV* is smaller than *LTD*, then *TV* is *SV* extended on the right by *LTD-M* `<space>`s.
- d) If *SD* is a datetime data type or an interval data type, then let *Y* be the shortest character string that conforms to the definition of `<literal>` in Subclause 5.3, “`<literal>`”, and such that the interpreted value of *Y* is *SV* and the interpreted precision of *Y* is the precision of *SD*. If *SV* is a negative interval, then `<sign>` shall be specified within `<unquoted interval string>` in the literal *Y*.

Case:

- i) If *Y* contains any `<SQL language character>` that is not in the character repertoire of *TD*, then an exception condition is raised: *data exception — invalid character value for cast*.
 - ii) If the length in characters *LY* of *Y* is equal to *LTD*, then *TV* is *Y*.
 - iii) If the length in characters *LY* of *Y* is less than *LTD*, then *TV* is *Y* extended on the right by *LTD-LY* `<space>`s.
 - iv) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- e) If *SD* is boolean, then

Case:

- i) If *SV* is *True* and *LTD* is not less than 4, then *TV* is 'TRUE' extended on the right by *LTD-4* `<space>`s.
- ii) If *SV* is *False* and *LTD* is not less than 5, then *TV* is 'FALSE' extended on the right by *LTD-5* `<space>`s.
- iii) Otherwise, an exception condition is raised: *data exception — invalid character value for cast*.

- 11) If *TD* is variable-length character string or large object character string, then let *MLTD* be the maximum length in characters of *TD*.

Case:

- a) If *SD* is exact numeric, then:

- i) Let *YP* be the shortest character string that conforms to the definition of `<exact numeric literal>` in Subclause 5.3, “`<literal>`”, whose scale is the same as the scale of *SD* and whose interpreted value is the absolute value of *SV*.

- ii) Case:

1) If *SV* is less than 0 (zero), then let *Y* be the result of ' - ' | | *YP*.

- 2) Otherwise, let *Y* be *YP*.
- iii) Case:
 - 1) If *Y* contains any <SQL language character> that is not in the character repertoire of *TD*, then an exception condition is raised: *data exception — invalid character value for cast*.
 - 2) If the length in characters *LY* of *Y* is less than or equal to *MLTD*, then *TV* is *Y*.
 - 3) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- b) If *SD* is approximate numeric, then
 - i) Let *YP* be a character string as follows.

Case:

 - 1) If *SV* equals 0 (zero), then *YP* is '0E0'.
 - 2) Otherwise, *YP* is the shortest character string that conforms to the definition of <approximate numeric literal> in Subclause 5.3, “<literal>”, whose interpreted value is equal to the absolute value of *SV* and whose <mantissa> consists of a single <digit> that is not '0', followed by a <period> and an <unsigned integer>.
 - ii) Case:
 - 1) If *SV* is less than 0 (zero), then let *Y* be the result of ' - ' || *YP*.
 - 2) Otherwise, let *Y* be *YP*.
 - iii) Case:
 - 1) If *Y* contains any <SQL language character> that is not in the character repertoire of *TD*, then an exception condition is raised: *data exception — invalid character value for cast*.
 - 2) If the length in characters *LY* of *Y* is less than or equal to *MLTD*, then *TV* is *Y*.
 - 3) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- c) If *SD* is fixed-length character string, variable-length character string, or large object character string, then

Case:

 - i) If the length in characters of *SV* is less than or equal to *MLTD*, then *TV* is *SV*.
 - ii) If the length in characters of *SV* is larger than *MLTD*, then *TV* is the first *MLTD* characters of *SV*. If any of the remaining characters of *SV* are non-<space> characters, then a completion condition is raised: *warning — string data, right truncation*.
- d) If *SD* is a datetime data type or an interval data type then let *Y* be the shortest character string that conforms to the definition of <literal> in Subclause 5.3, “<literal>”, and such that the interpreted value of *Y* is *SV* and the interpreted precision of *Y* is the precision of *SD*. If *SV* is a negative interval, then <sign> shall be specified within <unquoted interval string> in the literal *Y*.

Case:

 - i) If *Y* contains any <SQL language character> that is not in the character repertoire of *TD*, then an exception condition is raised: *data exception — invalid character value for cast*.

- ii) If the length in characters *LY* of *Y* is less than or equal to *MLTD*, then *TV* is *Y*.
- iii) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- e) If *SD* is boolean, then
 - Case:
 - i) If *SV* is *True* and *MLTD* is not less than 4, then *TV* is 'TRUE'.
 - ii) If *SV* is *False* and *MLTD* is not less than 5, then *TV* is 'FALSE'.
 - iii) Otherwise, an exception condition is raised: *data exception — invalid character value for cast*.

12) If *TD* and *SD* are binary string types, then

Case:

- a) If *TD* is fixed-length binary string, then let *LTD* be the length in octets of *TD*.
 - i) If the length in octets of *SV* is equal to *LTD*, then *TV* is *SV*.
 - ii) If the length in octets of *SV* is larger than *LTD*, then *TV* is the first *LTD* octets of *SV* and a completion condition is raised: *warning — string data, right truncation*.
 - iii) If the length in octets *M* of *SV* is smaller than *LTD*, then *TV* is *SV* extended on the right by *LTD-M* 'X'00's.
- b) Otherwise, let *MLTD* be the maximum length in octets of *TD*.

Case:

- i) If the length in octets of *SV* is less than or equal to *MLTD*, then *TV* is *SV*.
- ii) If the length in octets of *SV* is larger than *MLTD*, then *TV* is the first *MLTD* octets of *SV* and a completion condition is raised: *warning — string data, right truncation*.

13) If *TD* is the datetime data type DATE, then

Case:

- a) If *SD* is character string, then *SV* is replaced by

```
TRIM ( BOTH ' ' FROM VE )
```

Case:

- i) If the rules for <literal> or for <unquoted date string> in [Subclause 5.3](#), “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.
- ii) Otherwise, an exception condition is raised: *data exception — invalid datetime format*.
- b) If *SD* is the datetime data type DATE, then *TV* is *SV*.
- c) If *SD* is the datetime data type TIMESTAMP WITHOUT TIME ZONE, then *TV* is the year, month, and day <primary datetime field>s of *SV*.
- d) If *SD* is the datetime data type TIMESTAMP WITH TIME ZONE, then *TV* is computed by:

```
CAST ( CAST ( VE AS TIMESTAMP WITHOUT TIME ZONE ) AS DATE )
```

14) Let *STZD* be the current default time zone displacement of the SQL-session.

15) If *TD* is the datetime data type TIME WITHOUT TIME ZONE, then let *TSP* be the <time precision> of *TD*.

Case:

a) If *SD* is character string, then *SV* is replaced by:

```
TRIM ( BOTH ' ' FROM VE )
```

Case:

- i) If the rules for <literal> or for <unquoted time string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.
- ii) If the rules for <literal> or for <unquoted time string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type TIME(*TSP*) WITH TIME ZONE, then let *TV* be the value of:

```
CAST ( CAST ( VE AS TIME(TSP) WITH TIME ZONE ) AS  
TIME(TSP) WITHOUT TIME ZONE )
```

- iii) Otherwise, an exception condition is raised: *data exception — invalid datetime format*.
- b) If *SD* is TIME WITHOUT TIME ZONE, then *TV* is *SV*, with implementation-defined rounding or truncation if necessary.
- c) If *SD* is TIME WITH TIME ZONE, then let *SVUTC* be the UTC component of *SV* and let *SVTZ* be the time zone displacement of *SV*. *TV* is *SVUTC* + *SVTZ*, computed modulo 24 hours, with implementation-defined rounding or truncation if necessary.
- d) If *SD* is TIMESTAMP WITHOUT TIME ZONE, then *TV* is the hour, minute, and second <primary datetime field>s of *SV*, with implementation-defined rounding or truncation if necessary.
- e) If *SD* is TIMESTAMP WITH TIME ZONE, then *TV* is:

```
CAST ( CAST ( VE AS TIMESTAMP(TSP) WITHOUT TIME ZONE )  
AS TIME(TSP) WITHOUT TIME ZONE )
```

16) If *TD* is the datetime data type TIME WITH TIME ZONE, then let *TSP* be the <time precision> of *TD*.

Case:

a) If *SD* is character string, then *SV* is replaced by:

```
TRIM ( BOTH ' ' FROM VE )
```

Case:

- i) If the rules for <literal> or for <unquoted time string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.
- ii) If the rules for <literal> or for <unquoted time string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type TIME(*TSP*) WITHOUT TIME ZONE, then let *TV* be the value of:

```
CAST ( CAST ( VE AS TIME(TSP) WITHOUT TIME ZONE )
AS TIME(TSP) WITH TIME ZONE )
```

- iii) Otherwise, an exception condition is raised: *data exception — invalid datetime format*.
- b) If *SD* is TIME WITH TIME ZONE, then *TV* is *SV*, with implementation-defined rounding or truncation if necessary.
- c) If *SD* is TIME WITHOUT TIME ZONE, then the UTC component of *TV* is *SV* – *STZD*, computed modulo 24 hours, with implementation-defined rounding or truncation if necessary, and the time zone component of *TV* is *STZD*.
- d) If *SD* is TIMESTAMP WITH TIME ZONE, then the UTC component of *TV* is the hour, minute, and second <primary datetime field>s of *SV*, with implementation-defined rounding or truncation if necessary, and the time zone component of *TV* is the time zone displacement of *SV*.
- e) If *SD* is TIMESTAMP WITHOUT TIME ZONE, then *TV* is:

```
CAST ( CAST ( VE AS TIMESTAMP(TSP) WITH TIME ZONE )
AS TIME(TSP) WITH TIME ZONE )
```

- 17) If *TD* is the datetime data type TIMESTAMP WITHOUT TIME ZONE, then let *TSP* be the <timestamp precision> of *TD*.

Case:

- a) If *SD* is character string, then *SV* is replaced by:

```
TRIM ( BOTH ' ' FROM VE )
```

Case:

- i) If the rules for <literal> or for <unquoted timestamp string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.
- ii) If the rules for <literal> or for <unquoted timestamp string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type TIMESTAMP(*TSP*) WITH TIME ZONE, then let *TV* be the value of:

```
CAST ( CAST ( VE AS TIMESTAMP(TSP) WITH TIME ZONE )
AS TIMESTAMP(TSP) WITHOUT TIME ZONE )
```

- iii) Otherwise, an exception condition is raised: *data exception — invalid datetime format*.
- b) If *SD* is a date, then the <primary datetime field>s hour, minute, and second of *TV* are set to 0 (zero) and the <primary datetime field>s year, month, and day of *TV* are set to their respective values in *SV*.
- c) If *SD* is TIME WITHOUT TIME ZONE, then the <primary datetime field>s year, month, and day of *TV* are set to their respective values in an execution of CURRENT_DATE and the <primary datetime field>s hour, minute, and second of *TV* are set to their respective values in *SV*, with implementation-defined rounding or truncation if necessary.
- d) If *SD* is TIME WITH TIME ZONE, then *TV* is:

```
CAST ( CAST ( VE AS TIMESTAMP WITH TIME ZONE )
AS TIMESTAMP WITHOUT TIME ZONE )
```

- e) If *SD* is **TIMESTAMP WITHOUT TIME ZONE**, then *TV* is *SV*, with implementation-defined rounding or truncation if necessary.
- f) If *SD* is **TIMESTAMP WITH TIME ZONE**, then let *SVUTC* be the UTC component of *SV* and let *SVTZ* be the time zone displacement of *SV*. *TV* is *SVUTC* + *SVTZ*, with implementation-defined rounding or truncation if necessary.

18) If *TD* is the datetime data type **TIMESTAMP WITH TIME ZONE**, then let *TSP* be the <timestamp precision> of *TD*.

Case:

- a) If *SD* is character string, then *SV* is replaced by:

```
TRIM ( BOTH ' ' FROM VE )
```

Case:

- i) If the rules for <literal> or for <unquoted timestamp string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.
- ii) If the rules for <literal> or for <unquoted timestamp string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type **TIMESTAMP(*TSP*) WITHOUT TIME ZONE**, then let *TV* be the value of:

```
CAST ( CAST ( VE AS TIMESTAMP(TSP) WITHOUT TIME ZONE )
AS TIMESTAMP(TSP) WITH TIME ZONE )
```

- iii) Otherwise, an exception condition is raised: *data exception — invalid datetime format*.

- b) If *SD* is a date, then *TV* is:

```
CAST ( CAST ( VE AS TIMESTAMP(TSP) WITHOUT TIME ZONE )
AS TIMESTAMP(TSP) WITH TIME ZONE )
```

- c) If *SD* is **TIME WITHOUT TIME ZONE**, then *TV* is:

```
CAST ( CAST ( VE AS TIMESTAMP(TSP) WITHOUT TIME ZONE )
AS TIMESTAMP(TSP) WITH TIME ZONE )
```

- d) If *SD* is **TIME WITH TIME ZONE**, then the <primary datetime field>s year, month, and day of *TV* are set to their respective values in an execution of **CURRENT_DATE** and the <primary datetime field>s hour, minute, and second are set to their respective values in *SV*, with implementation-defined rounding or truncation if necessary. The time zone component of *TV* is set to the time zone component of *SV*.
- e) If *SD* is **TIMESTAMP WITHOUT TIME ZONE**, then the UTC component of *TV* is *SV* – *STZD*, with a time zone displacement of *STZD*.
- f) If *SD* is **TIMESTAMP WITH TIME ZONE**, then *TV* is *SV* with implementation-defined rounding or truncation, if necessary.

19) If *TD* is interval, then

Case:

- a) If *SD* is exact numeric, then

Case:

- i) If the representation of *SV* in the data type *TD* would result in the loss of leading significant digits, then an exception condition is raised: *data exception — interval field overflow*.
- ii) Let *NDESN* be the scale of the <exact numeric literal> formed by the <cast specification>

CAST (*SV* AS CHARACTER VARYING(*max*))

where *max* is the implementation-defined maximum precision of the CHARACTER VARYING type. If the number of digits of fractional seconds precision *NDFSP* of *TD* is less than *NDESN*, then it is implementation-defined whether *TV* is determined by rounding *SV* to *NDFSP* digits of precision or by truncating *SV* to *NDFSP* digits of precision, as specified in Subclause 4.4.2, “Characteristics of numbers”.

- iii) Otherwise, *TV* is that representation.
- b) If *SD* is character string, then *SV* is replaced by

TRIM (BOTH ' ' FROM *VE*)

Case:

- i) If the rules for <literal> or for <unquoted interval string> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.
 - ii) Otherwise, an exception condition is raised: *data exception — invalid interval format*.
- c) If *SD* is interval and *TD* and *SD* have the same interval precision, then *TV* is *SV*.
- d) If *SD* is interval and *TD* and *SD* have different interval precisions, then let *Q* be the least significant <primary datetime field> of *TD*.
- i) Let *Y* be the result of converting *SV* to a scalar in units *Q* according to the natural rules for intervals as defined in the Gregorian calendar (that is, there are 60 seconds in a minute, 60 minutes in an hour, 24 hours in a day, and 12 months in a year).
 - ii) Normalize *Y* to conform to the <interval qualifier> “*P* TO *Q*” of *TD* (again, observing the rules that there are 60 seconds in a minute, 60 minutes in an hour, 24 hours in a day, and 12 months in a year). Whether to truncate or round in the least significant field of the result is implementation-defined. If this would result in loss of precision of the leading datetime field of *Y*, then an exception condition is raised: *data exception — interval field overflow*.
 - iii) *TV* is the value of *Y*.

20) 09 If *TD* is boolean, then

Case:

- a) If *SD* is character string, then *SV* is replaced by

TRIM (BOTH ' ' FROM *VE*)

Case:

- i) If the rules for <literal> in Subclause 5.3, “<literal>”, can be applied to *SV* to determine a valid value of the data type *TD*, then let *TV* be that value.

- ii) Otherwise, an exception condition is raised: *data exception — invalid character value for cast*.
 - b) If *SD* is boolean, then *TV* is *SV*.
- 21) If the <cast specification> contains a <domain name> and that <domain name> refers to a domain that contains a <domain constraint> and if *TV* does not satisfy the <check constraint definition> simply contained in the <domain constraint>, then an exception condition is raised: *integrity constraint violation*.
- 22) The result of *CS* is *TV*.

Conformance Rules

- 1) Without Feature T042, “Extended LOB data type support”, conforming SQL language shall not contain a <cast operand> whose declared type is BINARY LARGE OBJECT or CHARACTER LARGE OBJECT.
- 2) Without Feature F421, “National character”, conforming SQL language shall not contain a <cast operand> whose declared type is NATIONAL CHARACTER LARGE OBJECT.
- 3) Without Feature T042, “Extended LOB data type support”, conforming SQL language shall not contain a <cast operand> whose declared type is NATIONAL CHARACTER LARGE OBJECT.
- 4) Without Feature S043, “Enhanced reference types”, in conforming SQL language, if the declared type of <cast operand> is a reference type, then <cast target> shall contain a <data type> that is a reference type.

6.13 <next value expression>

This Subclause is modified by Subclause 6.3, “<next value expression>”, in ISO/IEC 9075-4.

Function

Return the next value of a sequence generator.

Format

```
<next value expression> ::=
  NEXT VALUE FOR <sequence generator name>
```

Syntax Rules

- 1) A <next value expression> shall be directly contained in one of the following:
 - a) A <select list> simply contained in a <query specification> that constitutes a <query expression> that is immediately contained in one of the following:
 - i) A <cursor specification>.
 - ii) A <table subquery> simply contained in an <as subquery clause> in a <table definition>.
 - iii) A <from subquery>.
 - iv) A <select statement: single row>.
 - b) A <select list> simply contained in a <query specification> that is immediately contained in a <dynamic single row select statement>.
 - c) A <from constructor>.
 - d) A <merge insert value list>.
 - e) 04 An <update source>.
- 2) <next value expression> shall not be contained in a <case expression>, a <search condition>, an <order by clause>, an <aggregate function>, a <window function>, a grouped query, or in a <query specification> that simply contains the <set quantifier> DISTINCT.
- 3) The declared type of <next value expression> is the data type described by the data type descriptor included in the sequence generator descriptor identified by <sequence generator name>.

Access Rules

- 1) Case:
 - a) If <next value expression> is contained in a <schema definition>, then the applicable privileges for the <authorization identifier> that owns the containing schema shall include USAGE privilege on the sequence generator identified by <sequence generator name>.

- b) Otherwise, the current privileges shall include USAGE privilege on the sequence generator identified by <sequence generator name>.

General Rules

- 1) If <next value expression> *NVE* is specified, then let *SEQ* be the sequence generator descriptor identified by the <sequence generator name> contained in *NVE*.

Case:

- a) If *NVE* is directly contained in a <query specification> *QS*, then the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”, are applied once per row in the result of *QS* with *SEQ* as *SEQUENCE*. The result of each evaluation of *NVE* for a given row is the *RESULT* returned by the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”.
- b) If *NVE* is directly contained in a <contextually typed table value constructor> *TVC*, then the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”, are applied once per <contextually typed row value expression> contained in *TVC* with *SEQ* as *SEQUENCE*. The result of each evaluation of *NVE* for a given <row value expression> is the *RESULT* returned by the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”.
- c) If *NVE* is directly contained in an <update source>, then the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”, are applied once per row to be updated by the <update statement: searched> or <update statement: positioned> with *SEQ* as *SEQUENCE*. The result of each evaluation of *NVE* for a given row is the *RESULT* returned by the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”.

Conformance Rules

- 1) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <next value expression>.

6.14 <field reference>

Function

Reference a field of a row value.

Format

```
<field reference> ::=  
    <value expression primary> <period> <field name>
```

Syntax Rules

- 1) Let *FR* be the <field reference>, let *VEP* be the <value expression primary> immediately contained in *FR*, and let *FN* be the <field name> immediately contained in *FR*.
- 2) The declared type of *VEP* shall be a row type. Let *RT* be that row type.
- 3) *FR* is a *field reference*.
- 4) *FN* shall unambiguously reference a field of *RT*. Let *F* be that field.
- 5) The declared type of *FR* is the declared type of *F*.

Access Rules

None.

General Rules

- 1) Let *VR* be the value of *VEP*.
- 2) Case:
 - a) If *VR* is the null value, then the value of *FR* is the null value.
 - b) Otherwise, the value of *FR* is the value of the field *F* of *VR*.

Conformance Rules

- 1) Without Feature T051, “Row types”, conforming SQL language shall not contain a <field reference>.

6.15 <subtype treatment>

Function

Modify the declared type of an expression.

Format

```
<subtype treatment> ::=  
    TREAT <left paren> <subtype operand> AS <target subtype> <right paren>  
  
<subtype operand> ::=  
    <value expression>  
  
<target subtype> ::=  
    <path-resolved user-defined type name>  
    | <reference type>
```

Syntax Rules

- 1) The declared type *VT* of the <value expression> shall be a structured type or a reference type.
- 2) Case:
 - a) If *VT* is a structured type, then:
 - i) <target subtype> shall specify a <path-resolved user-defined type name>.
 - ii) Let *DT* be the structured type identified by the <user-defined type name> simply contained in <path-resolved user-defined type name>.
 - b) Otherwise:
 - i) <target subtype> shall specify a <reference type>.
 - ii) Let *DT* be the reference type identified by <reference type>.
- 3) *VT* shall be a supertype of *DT*.
- 4) The declared type of the result of the <subtype treatment> is *DT*.

Access Rules

None.

General Rules

- 1) Let *V* be the value of the <value expression>.
- 2) Case:
 - a) If *V* is the null value, then the value of the <subtype treatment> is the null value.

6.15 <subtype treatment>

b) Otherwise:

- i) If the most specific type of *V* is not a subtype of *DT*, then an exception condition is raised: *invalid target type specification*.

NOTE 126 — “most specific type” is defined in [Subclause 4.7.5, “Subtypes and supertypes”](#).

- ii) The value of the <subtype treatment> is *V*.

Conformance Rules

- 1) Without Feature S161, “Subtype treatment”, conforming SQL Language shall not contain a <subtype treatment>.
- 2) Without Feature S162, “Subtype treatment for references”, conforming SQL language shall not contain a <target subtype> that contains a <reference type>.

6.16 <method invocation>

This Subclause is modified by *Subclause 6.1, “<method invocation>”, in ISO/IEC 9075-13.*

Function

Reference an SQL-invoked method of a user-defined type value.

Format

```
<method invocation> ::=
    <direct invocation>
  | <generalized invocation>

<direct invocation> ::=
    <value expression primary> <period> <method name> [ <SQL argument list> ]

<generalized invocation> ::=
    <left paren> <value expression primary> AS <data type> <right paren>
    <period> <method name> [ <SQL argument list> ]

<method selection> ::=
    <routine invocation>

<constructor method selection> ::=
    <routine invocation>
```

Syntax Rules

- 1) Let *OR* be the <method invocation>, let *VEP* be the <value expression primary> immediately contained in the <direct invocation> or <generalized invocation> of *OR*, and let *MN* be the <method name> immediately contained in *OR*.
- 2) 13 The declared type of *VEP* shall be a user-defined type. Let *UDT* be that user-defined type.
- 3) Case:
 - a) If <SQL argument list> is specified, then let *AL* be:

$$, A_1, \dots, A_n$$
 where A_i , $1 \text{ (one)} \leq i \leq n$, are the <SQL argument>s immediately contained in <SQL argument list>, taken in order of their ordinal position in <SQL argument list>.
 - b) Otherwise, let *AL* be a zero-length string.
- 4) Case:
 - a) If <method invocation> is immediately contained in <new invocation>, then let *TP* be an SQL-path containing the <schema name> of the schema that includes the descriptor of *UDT*.
 - b) Otherwise, let *TP* be an SQL-path, arbitrarily defined, containing the <schema name> of every schema that includes a descriptor of a supertype or subtype of *UDT*.

5) Case:

- a) If <generalized invocation> is specified, then let *DT* be the <data type> simply contained in the <generalized invocation>. Let *RI* be the following <method selection>:

MN (*VEP AS DT AL*)

b) Otherwise,

Case:

- i) If <method invocation> is immediately contained in <new invocation>, then let *RI* be the <constructor method selection>:

MN (*VEP AL*)

- ii) Otherwise, let *RI* be the following <method selection>:

MN (*VEP AL*)

- 6) The Syntax Rules of Subclause 10.4, “<routine invocation>”, are applied with *RI* as *ROUTINE INVOCATION* and *TP* as *SQLPATH*, yielding subject routine *SR* and static SQL argument list *SAL*.

Access Rules

None.

General Rules

- 1) The General Rules of Subclause 10.4, “<routine invocation>”, are applied with *SR* as *SUBJECT ROUTINE* and *SAL* as *STATIC SQL ARG LIST*, yielding value *V* that is the result of the <routine invocation>.
- 2) The value of <method invocation> is *V*.

Conformance Rules

- 1) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <method invocation>.

6.17 <static method invocation>

Function

Invoke a static method.

Format

```
<static method invocation> ::=  
  <path-resolved user-defined type name> <double colon> <method name>  
    [ <SQL argument list> ]  
  
<static method selection> ::=  
  <routine invocation>
```

Syntax Rules

- 1) Let *TN* be the <user-defined type name> immediately contained in <path-resolved user-defined type name> and let *T* be the user-defined type identified by *TN*.
- 2) Let *MN* be the <method name> immediately contained in <static method invocation>.
- 3) Case:
 - a) If <SQL argument list> is specified, then let *AL* be that <SQL argument list>.
 - b) Otherwise, let *AL* be <left paren> <right paren>.
- 4) Let *TP* be an SQL-path containing only the <schema name> of every schema that includes a descriptor of a supertype of *T*.
- 5) Let *RI* be the following <routine invocation>:

MN AL
- 6) Let *SMS* be the following <static method selection>:

RI
- 7) The Syntax Rules of Subclause 10.4, “<routine invocation>”, are applied with *RI* as *ROUTINE INVOCATION*, *TP* as *SQLPATH*, and *T* as *UDT*, yielding subject routine *SR* and static SQL argument list *SAL*.

Access Rules

None.

General Rules

- 1) The General Rules of Subclause 10.4, “<routine invocation>”, are applied with *SR* as *SUBJECT ROUTINE* and *SAL* as *STATIC SQL ARG LIST*, yielding a value *V* that is the result of the <routine invocation>.

- 2) The value of <static method invocation> is V.

Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <static method invocation>.

6.18 <new specification>

This Subclause is modified by *Subclause 6.2, “<new specification>”, in ISO/IEC 9075-13.*

Function

Invoke a method on a newly-constructed value of a structured type.

Format

```
<new specification> ::=  
  NEW <path-resolved user-defined type name> <SQL argument list>  
  
<new invocation> ::=  
  <method invocation>  
  | <routine invocation>
```

Syntax Rules

- 1) Let *UDTN* be the <path-resolved user-defined type name> immediately contained in the <new specification>. Let *MN* be the <qualified identifier> immediately contained in *UDTN*.
- 2) Let *UDT* be the user-defined type identified by *UDTN*. *UDT* shall be instantiable. Let *SN* be the implicit or explicit <schema name> of *UDTN*. Let *S* be the schema identified by *SN*. Let *RN* be *NS.MN*.
- 3) Case:
 - a) If the <new specification> is of the form

NEW UDTN()

then

Case:

- i) If *S* does not include the descriptor of an SQL-invoked constructor method whose method name is equivalent to *MN* and whose unaugmented parameter list is empty, then the <new specification> is equivalent to the <new invocation>

RN()

- ii) Otherwise, the <new specification> is equivalent to the <new invocation>

RN().MN()

- b) Otherwise, the <new specification>

NEW UDTN(a₁, a₂, ..., a_n)

is equivalent to the <new invocation>

RN().MN(a₁, a₂, ..., a_n)

Access Rules

None.

NOTE 127 — The applicable privileges or current privileges (as appropriate) include EXECUTE privilege on the constructor function, and also on the indicated constructor method, according to the Syntax Rules of [Subclause 10.4](#), “<routine invocation>”.

General Rules

None.

Conformance Rules

- 1) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <new specification>.

6.19 <attribute or method reference>

Function

Return a value acquired by accessing a column of the row identified by a value of a reference type or by invoking an SQL-invoked method.

Format

```
<attribute or method reference> ::=  
  <value expression primary> <dereference operator> <qualified identifier>  
    [ <SQL argument list> ]  
  
<dereference operator> ::=  
  <right arrow>
```

Syntax Rules

- 1) The declared type of the <value expression primary> *VEP* shall be a reference type and the scope included in its reference type descriptor shall not be empty. Let *RT* be the referenced type of *VEP*.
- 2) Let *QI* be the <qualified identifier>. If <SQL argument list> is specified, then let *SAL* be <SQL argument list>; otherwise, let *SAL* be a zero-length string.
- 3) Case:
 - a) If *QI* is equivalent to the attribute name of an attribute of *RT* and *SAL* is a zero-length string, then <attribute or method reference> is effectively replaced by a <dereference operation> *AMR* of the form:
$$VEP \rightarrow QI$$
 - b) Otherwise, <attribute or method reference> is effectively replaced by a <method reference> *AMR* of the form:
$$VEP \rightarrow QI \text{ } SAL$$
- 4) The declared type of <attribute or method reference> is the declared type of *AMR*.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain an <attribute or method reference>.

6.20 <dereference operation>

Function

Access a column of the row identified by a value of a reference type.

Format

```
<dereference operation> ::=  
  <reference value expression> <dereference operator> <attribute name>
```

Syntax Rules

- 1) Let *RVE* be the <reference value expression>. The reference type descriptor of *RVE* shall include a scope. Let *RT* be the referenced type of *RVE*.
- 2) Let *AN* be the <attribute name>. *AN* shall identify an attribute *AT* of *RT*.
- 3) The declared type of the <dereference operation> is the declared type of *AT*.
- 4) Let *S* be the name of the referenceable table in the scope of the reference type of *RVE*.
- 5) Let *OID* be the name of the self-referencing column of *S*.
- 6) <dereference operation> is equivalent to a <scalar subquery> of the form:

```
( SELECT AN  
  FROM S  
  WHERE S.OID = RVE )
```

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <dereference operation>.

6.21 <method reference>

Function

Return a value acquired from invoking an SQL-invoked routine that is a method.

Format

```
<method reference> ::=  
  <value expression primary> <dereference operator> <method name> <SQL argument list>
```

Syntax Rules

- 1) The data type of the <value expression primary> *VEP* shall be a reference type and the scope included in its reference type descriptor shall not be empty.
- 2) Let *MN* be the method name. Let *MRAL* be the <SQL argument list>.
- 3) The Syntax Rules of [Subclause 6.16](#), “<method invocation>”, are evaluated for the <method invocation>:

DEREF (VEP) . MN MRAL

yielding subject routine *SR* and static SQL argument list *SAL*.

- 4) The data type of <method reference> is the data type of the expression:

DEREF (VEP) . MN MRAL

Access Rules

- 1) Let *SCOPE* be the table that is the scope of *VEP*.

Case:

- a) If <method reference> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include the table/method privilege for table *SCOPE* and method *SR*.
- b) Otherwise, the current privileges shall include the table/method privilege for table *SCOPE* and method *SR*.

General Rules

- 1) The General Rules of [Subclause 6.16](#), “<method invocation>”, are applied with *SR* and *SAL* as the subject routine and SQL argument list, respectively, yielding a value *V* that is the result of the <routine invocation>.
- 2) The value of <method reference> is *V*.

Conformance Rules

- 1) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <method reference>.

6.22 <reference resolution>

Function

Obtain the value referenced by a reference value.

Format

```
<reference resolution> ::=  
  DEREf <left paren> <reference value expression> <right paren>
```

Syntax Rules

- 1) Let *RR* be the <reference resolution> and let *RVE* be the <reference value expression>. The reference type descriptor of *RVE* shall include a scope.
- 2) The declared type of *RR* is the structured type that is referenced by the declared type of *RVE*.
- 3) Let *SCOPE* be the table identified by the table name included in the reference type descriptor of *RVE*. *SCOPE* is the scoped table of *RR*.

NOTE 128 — The term “scoped table” is defined in [Subclause 4.9, “Reference types”](#).

- 4) Let *m* be the number of subtables of *SCOPE*. Let *S_i*, 1 (one) ≤ *i* ≤ *m*, be the subtables, arbitrarily ordered, of *SCOPE*.
- 5) For each *S_i*, 1 (one) ≤ *i* ≤ *m*, let *STN_i* be the name included in the descriptor of *S_i* of the structured type *ST_i* associated with *S_i*, let *REFCOL_i* be the self-referencing column of *S_i*, let *N_i* be the number of attributes of *ST_i*, and let *A_{i,j}*, 1 (one) ≤ *j* ≤ *N_i*, be the names of the attributes of *ST_i*, therefore also the names of the columns of *S_i*.

Access Rules

- 1) Case:
 - a) If <reference resolution> is contained in a <schema definition>, then the applicable privileges for the <authorization identifier> that owns the containing schema shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *SCOPE*.
 - b) Otherwise, the current privileges shall includes SELECT WITH HIERARCHY OPTION on at least one supertable of *SCOPE*.

General Rules

- 1) The value of <reference resolution> is the value of:

```
(  
  SELECT A1,1 ( ... A1,N1  
    ( STN1( ), A1,N1 ), ... A1,1 )
```



```

FROM ONLY  $S_1$ 
WHERE  $S_1.REFCOL_1 = RVE$ 
UNION
SELECT  $A_{2,1}$  ( ...  $A_{2,N_2}$ 
              (  $STN_2()$ ,  $A_{2,N_2}$  ), ...  $A_{2,1}$  )
FROM ONLY  $S_2$ 
WHERE  $S_2.REFCOL_2 = RVE$ 
UNION
...
UNION
SELECT  $A_{m,1}$  ( ...  $A_{m,N_m}$ 
              (  $STN_m()$ ,  $A_{m,N_m}$  ), ...  $A_{m,1}$  )
FROM ONLY  $S_m$ 
WHERE  $S_m.REFCOL_m = RVE$ 
)

```

NOTE 129 — The evaluation of this General Rule is effectively performed without further Access Rule checking.

Conformance Rules

- 1) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <reference resolution>.

6.23 <array element reference>

Function

Return an element of an array.

Format

```
<array element reference> ::=
  <array value expression>
    <left bracket or trigraph> <numeric value expression> <right bracket or trigraph>
```

Syntax Rules

- 1) The declared type of an <array element reference> is the element type of the specified <array value expression>.
- 2) The declared type of <numeric value expression> shall be exact numeric with scale 0 (zero).

Access Rules

None.

General Rules

- 1) If the value of <array value expression> or <numeric value expression> is the null value, then the result of <array element reference> is the null value.
- 2) Let the value of <numeric value expression> be *i*.

Case:

- a) If *i* is greater than zero and less than or equal to the cardinality of <array value expression>, then the result of <array element reference> is the value of the *i*-th element of the value of <array value expression>.
- b) Otherwise, an exception condition is raised: *data exception — array element error*.

Conformance Rules

- 1) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array element reference>.

6.24 <multiset element reference>

Function

Return the sole element of a multiset of one element.

Format

```
<multiset element reference> ::=
  ELEMENT <left paren> <multiset value expression> <right paren>
```

Syntax Rules

- 1) Let *MVE* be the <multiset value expression>. The <multiset element reference> is equivalent to the <scalar subquery>

```
( SELECT M.E
  FROM UNNEST (MVE) AS M(E) )
```

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset element reference>.

6.25 <value expression>

This Subclause is modified by Subclause 6.3, “<value expression>”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 6.7, “<value expression>”, in ISO/IEC 9075-14.

Function

Specify a value.

Format

```
<value expression> ::=
    <common value expression>
  | <boolean value expression>
  | <row value expression>

09 14 <common value expression> ::=
    <numeric value expression>
  | <string value expression>
  | <datetime value expression>
  | <interval value expression>
  | <user-defined type value expression>
  | <reference value expression>
  | <collection value expression>

<user-defined type value expression> ::=
    <value expression primary>

<reference value expression> ::=
    <value expression primary>

<collection value expression> ::=
    <array value expression>
  | <multiset value expression>
```

Syntax Rules

- 1) The declared type of a <value expression> is the declared type of the simply contained <common value expression>, <boolean value expression>, or <row value expression>.
- 2) 09 14 The declared type of a <common value expression> is the declared type of the <numeric value expression>, <string value expression>, <datetime value expression>, <interval value expression>, <user-defined type value expression>, <collection value expression>, or <reference value expression>, respectively.
- 3) The declared type of a <user-defined type value expression> is the declared type of the immediately contained <value expression primary>, which shall be a user-defined type.
- 4) The declared type of a <reference value expression> is the declared type of the immediately contained <value expression primary>, which shall be a reference type.
- 5) The declared type of a <collection value expression> is the declared type of the immediately contained <array value expression> or <multiset value expression>.

- 6) Let *C* be some column. Let *VE* be the <value expression>. *C* is an underlying column of *VE* if and only if *C* is identified by some column reference contained in *VE*.
- 7) A <value expression> or <nonparenthesized value expression primary> is *possibly non-deterministic* if it generally contains any of the following:
- a) A <datetime value function>.
 - b) A <next value expression>.
 - c) A <cast specification> that either is, or recursively implies through the execution of the General Rules of Subclause 6.12, “<cast specification>”, one of the following:
 - i) A <cast specification> whose result type is datetime with time zone and whose <cast operand> has declared type that is not datetime with time zone.
 - ii) 14A <cast specification> whose result type is an array type and whose <cast operand> has a declared type that is a multiset type.
 - d) An <array value constructor by query>.
 - e) A <datetime factor> that simply contains a <datetime primary> whose declared type is datetime without time zone and that simply contains an explicit <time zone>.
 - f) An <interval value expression> that computes the difference of a <datetime value expression> and a <datetime term>, such that the declared type of one operand is datetime with time zone and the other operand is datetime without time zone.
 - g) A <comparison predicate>, <overlaps predicate>, or <distinct predicate> simply containing <row value predicand>s *RVP1* and *RVP2* such that the declared types of *RVP1* and *RVP2* have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.

NOTE 130 — This includes <between predicate> because of a syntactic transformation to <comparison predicate>.

- h) A <quantified comparison predicate> or a <match predicate> simply containing a <row value predicand> *RVP* and a <table subquery> *TS* such that the declared types of *RVP* and *TS* have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.

NOTE 131 — This includes <in predicate> because of a syntactic transformation to <quantified comparison predicate>.

- i) A <member predicate> simply containing a <row value predicand> *RVP* and a <multiset value expression> *MVP* such that the declared type of the only field *F* of *RVP* and the element type of *MVP* have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.
- j) A <submultiset predicate> simply containing a <row value predicand> *RVP* and a <multiset value expression> *MVP* such that the declared type of the only field *F* of *RVP* and the declared type of *MVP* have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.
- k) A <multiset value expression> that specifies or implies MULTiset UNION, MULTiset EXCEPT, or MULTiset INTERSECT such that the element types of the operands have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.
- l) A <value specification> that is CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, CURRENT_SCHEMA, or CURRENT_PATH.

- m) A <routine invocation> whose subject routine is an SQL-invoked routine that is possibly non-deterministic.
- n) 14 An <aggregate function> that specifies MIN or MAX and that simply contains a <value expression> whose declared type is based on a character string type, user-defined type, or datetime with time zone type.
- o) An <aggregate function> that specifies INTERSECTION and that simply contains a <value expression> whose declared element type is based on a character string type, a user-defined type, or a datetime type with time zone.
- p) An <array aggregate function>.
- q) A <multiset value expression> that specifies MULTiset UNION DISTINCT, MULTiset EXCEPT, or MULTiset INTERSECT and whose result type's declared element type is based on character string type, a user-defined type, or a datetime type with time zone.
- r) A <multiset set function> whose declared element type is based on a character string type, a user-defined type, or a datetime type with time zone.
- s) A <window function> that specifies ROW_NUMBER or whose associated <window specification> specifies ROWS.
- t) 14 A <query specification> or <query expression> that is possibly non-deterministic.

NOTE 132 — The notion of one data type being based on another data type is defined in Subclause 4.1, “Data types”.

Access Rules

None.

General Rules

- 1) The value of a <value expression> is the value of the simply contained <common value expression>, <boolean value expression>, or <row value expression>.
- 2) 14 The value of a <common value expression> is the value of the immediately contained <numeric value expression>, <string value expression>, <datetime value expression>, <interval value expression>, <user-defined type value expression>, <collection value expression>, or <reference value expression>.
- 3) When a <value expression> *V* is evaluated for a row *R* of a table, each reference to a column of that table by a column reference *CR* directly contained in *V* is the value of that column in that row.
- 4) The value of a <collection value expression> is the value of its immediately contained <array value expression> or <multiset value expression>.
- 5) The value of a <reference value expression> *RVE* is the value of the <value expression primary> immediately contained in *RVE*.

Conformance Rules

- 1) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <value expression> that is a <boolean value expression>.

- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <reference value expression>.

6.26 <numeric value expression>

Function

Specify a numeric value.

Format

```

<numeric value expression> ::=
    <term>
  | <numeric value expression> <plus sign> <term>
  | <numeric value expression> <minus sign> <term>

<term> ::=
    <factor>
  | <term> <asterisk> <factor>
  | <term> <solidus> <factor>

<factor> ::=
    [ <sign> ] <numeric primary>

<numeric primary> ::=
    <value expression primary>
  | <numeric value function>

```

Syntax Rules

- 1) If the declared type of both operands of a dyadic arithmetic operator is exact numeric, then the declared type of the result is an implementation-defined exact numeric type, with precision and scale determined as follows:
 - a) Let *S1* and *S2* be the scale of the first and second operands respectively.
 - b) The precision of the result of addition and subtraction is implementation-defined, and the scale is the maximum of *S1* and *S2*.
 - c) The precision of the result of multiplication is implementation-defined, and the scale is *S1* + *S2*.
 - d) The precision and scale of the result of division are implementation-defined.
- 2) If the declared type of either operand of a dyadic arithmetic operator is approximate numeric, then the declared type of the result is an implementation-defined approximate numeric type.
- 3) The declared type of a <factor> is that of the immediately contained <numeric primary>.
- 4) The declared type of a <numeric primary> shall be numeric.
- 5) If a <numeric value expression> immediately contains a <minus sign> *NMS* and immediately contains a <term> that immediately contains a <factor> that immediately contains a <sign> that is a <minus sign> *FMS*, then there shall be a <separator> between *NMS* and *FMS*.

Access Rules

None.

General Rules

- 1) If the value of any <numeric primary> simply contained in a <numeric value expression> is the null value, then the result of the <numeric value expression> is the null value.
- 2) If the <numeric value expression> contains only a <numeric primary>, then the result of the <numeric value expression> is the value of the specified <numeric primary>.
- 3) The monadic arithmetic operators <plus sign> and <minus sign> (+ and −, respectively) specify monadic plus and monadic minus, respectively. Monadic plus does not change its operand. Monadic minus reverses the sign of its operand.
- 4) The dyadic arithmetic operators <plus sign>, <minus sign>, <asterisk>, and <solidus> (+, −, *, and /, respectively) specify addition, subtraction, multiplication, and division, respectively. If the value of a divisor is zero, then an exception condition is raised: *data exception — division by zero*.
- 5) If the most specific type of the result of an arithmetic operation is exact numeric, then

Case:

- a) If the operator is not division and the mathematical result of the operation is not exactly representable with the precision and scale of the result data type, then an exception condition is raised: *data exception — numeric value out of range*.
 - b) If the operator is division and the approximate mathematical result of the operation represented with the precision and scale of the result data type loses one or more leading significant digits after rounding or truncating if necessary, then an exception condition is raised: *data exception — numeric value out of range*. The choice of whether to round or truncate is implementation-defined.
- 6) If the most specific type of the result of an arithmetic operation is approximate numeric and the exponent of the approximate mathematical result of the operation is not within the implementation-defined exponent range for the result data type, then an exception condition is raised: *data exception — numeric value out of range*.

Conformance Rules

None.

6.27 <numeric value function>

Function

Specify a function yielding a value of type numeric.

Format

```

<numeric value function> ::=
    <position expression>
  | <regex occurrences function>
  | <regex position expression>
  | <extract expression>
  | <length expression>
  | <cardinality expression>
  | <absolute value expression>
  | <modulus expression>
  | <natural logarithm>
  | <exponential function>
  | <power function>
  | <square root>
  | <floor function>
  | <ceiling function>
  | <width bucket function>

<position expression> ::=
    <character position expression>
  | <binary position expression>

<regex occurrences function> ::=
    OCCURRENCES_REGEX <left paren>
      <XQuery pattern> [ FLAG <XQuery option flag> ]
    IN <regex subject string>
      [ FROM <start position> ]
      [ USING <char length units> ]
    <right paren>

<XQuery pattern> ::=
    <character value expression>

<XQuery option flag> ::=
    <character value expression>

<regex subject string> ::=
    <character value expression>

<regex position expression> ::=
    POSITION_REGEX <left paren>
      [ <regex position start or after> ]
    <XQuery pattern> [ FLAG <XQuery option flag> ]
    IN <regex subject string>
      [ FROM <start position> ]
      [ USING <char length units> ]
      [ OCCURRENCE <regex occurrence> ]
      [ GROUP <regex capture group> ]

```

```
<right paren>

<regex position start or after> ::=
    START
    | AFTER

<regex occurrence> ::=
    <numeric value expression>

<regex capture group> ::=
    <numeric value expression>

<character position expression> ::=
    POSITION <left paren> <character value expression 1> IN <character value expression 2>
    [ USING <char length units> ] <right paren>

<character value expression 1> ::=
    <character value expression>

<character value expression 2> ::=
    <character value expression>

<binary position expression> ::=
    POSITION <left paren> <binary value expression> IN <binary value expression> <right paren>

<length expression> ::=
    <char length expression>
    | <octet length expression>

<char length expression> ::=
    { CHAR_LENGTH | CHARACTER_LENGTH } <left paren> <character value expression>
    [ USING <char length units> ] <right paren>

<octet length expression> ::=
    OCTET_LENGTH <left paren> <string value expression> <right paren>

<extract expression> ::=
    EXTRACT <left paren> <extract field> FROM <extract source> <right paren>

<extract field> ::=
    <primary datetime field>
    | <time zone field>

<time zone field> ::=
    TIMEZONE_HOUR
    | TIMEZONE_MINUTE

<extract source> ::=
    <datetime value expression>
    | <interval value expression>

<cardinality expression> ::=
    CARDINALITY <left paren> <collection value expression> <right paren>

<absolute value expression> ::=
    ABS <left paren> <numeric value expression> <right paren>

<modulus expression> ::=
    MOD <left paren> <numeric value expression dividend> <comma>
    <numeric value expression divisor> <right paren>
```

6.27 <numeric value function>

```

<numeric value expression dividend> ::=
    <numeric value expression>

<numeric value expression divisor> ::=
    <numeric value expression>

<natural logarithm> ::=
    LN <left paren> <numeric value expression> <right paren>

<exponential function> ::=
    EXP <left paren> <numeric value expression> <right paren>

<power function> ::=
    POWER <left paren> <numeric value expression base> <comma>
        <numeric value expression exponent> <right paren>

<numeric value expression base> ::=
    <numeric value expression>

<numeric value expression exponent> ::=
    <numeric value expression>

<square root> ::=
    SQRT <left paren> <numeric value expression> <right paren>

<floor function> ::=
    FLOOR <left paren> <numeric value expression> <right paren>

<ceiling function> ::=
    { CEIL | CEILING } <left paren> <numeric value expression> <right paren>

<width bucket function> ::=
    WIDTH_BUCKET <left paren> <width bucket operand> <comma> <width bucket bound 1> <comma>
        <width bucket bound 2> <comma> <width bucket count> <right paren>

<width bucket operand> ::=
    <numeric value expression>

<width bucket bound 1> ::=
    <numeric value expression>

<width bucket bound 2> ::=
    <numeric value expression>

<width bucket count> ::=
    <numeric value expression>

```

Syntax Rules

- 1) If <position expression>, <regex occurrences function>, or <regex position expression> is specified, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
- 2) If <character position expression> is specified, then <character value expression 1> and <character value expression 2> shall be comparable.
- 3) If <regex occurrences function> is specified, then:
 - a) <XQuery pattern> and <regex subject string> shall be comparable.

- b) The declared type of <start position> shall be exact numeric with scale 0 (zero).
 - c) If <start position> is not specified, then 1 (one) is implicit.
 - d) Case:
 - i) If <char length units> is specified, then the character repertoire of the <regex subject string> shall be UCS.
 - ii) Otherwise, CHARACTERS is implicit.
 - e) If <XQuery option flag> is not specified, then the zero-length string is implicit.
- 4) If <regex position expression> is specified, then:
- a) <XQuery pattern> and <regex subject string> shall be comparable.
 - b) If <regex position start or after> is not specified, then START is implicit.
 - c) The declared type of <start position>, <regex occurrence>, and <regex capture group> shall be exact numeric with scale 0 (zero).
 - d) If <start position> is not specified, then 1 (one) is implicit.
 - e) Case:
 - i) If <char length units> is specified, then the character repertoire of the <regex subject string> shall be UCS.
 - ii) Otherwise, CHARACTERS is implicit.
 - f) If <regex occurrence> is not specified, then 1 (one) is implicit.
 - g) If <regex capture group> is not specified, then 0 (zero) is implicit.
 - h) If <XQuery option flag> is not specified, then the zero-length string is implicit.
- 5) Case:
- a) If the character encoding form of <string value expression> is not UTF8, UTF16, or UTF32, then <char length units> shall not be specified.
 - b) Otherwise, if <char length units> is not specified, then CHARACTERS is implicit.
- 6) If <extract expression> is specified, then
- Case:
- a) If <extract field> is a <primary datetime field>, then it shall identify a <primary datetime field> of the <interval value expression> or <datetime value expression> immediately contained in <extract source>.
 - b) If <extract field> is a <time zone field>, then the declared type of the <extract source> shall be TIME WITH TIME ZONE or TIMESTAMP WITH TIME ZONE.
- 7) If <extract expression> is specified, then
- Case:
- a) If <primary datetime field> does not specify SECOND, then the declared type of the result is an implementation-defined exact numeric with scale 0 (zero).

6.27 <numeric value function>

- b) Otherwise, the declared type of the result is an implementation-defined exact numeric type with scale not less than the specified or implied <time fractional seconds precision> or <interval fractional seconds precision>, as appropriate, of the SECOND <primary datetime field> of the <extract source>.
- 8) If a <length expression> is specified, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
- 9) If <cardinality expression> is specified, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
- 10) If <absolute value expression> is specified, then the declared type of the result is the declared type of the immediately contained <numeric value expression>.
- 11) If <modulus expression> is specified, then the declared type of each <numeric value expression> shall be exact numeric with scale 0 (zero). The declared type of the result is the declared type of the immediately contained <numeric value expression divisor>.
- 12) The declared type of the result of <natural logarithm> is an implementation-defined approximate numeric type.
- 13) The declared type of the result of <exponential function> is an implementation-defined approximate numeric type.
- 14) The declared type of the result of <power function> is an implementation-defined approximate numeric type.
- 15) If <square root> is specified, then let *NVE* be the simply contained <numeric value expression>. The <square root> is equivalent to

`POWER (NVE, 0.5)`

- 16) If <floor function> or <ceiling function> is specified, then
Case:
 - a) If the declared type of the simply contained <numeric value expression> *NVE* is exact numeric, then the declared type of the result is exact numeric with implementation-defined precision, with the radix of *NVE*, and with scale 0 (zero).
 - b) Otherwise, the declared type of the result is approximate numeric with implementation-defined precision.
- 17) If <width bucket function> is specified, then the declared type of <width bucket count> shall be exact numeric with scale 0 (zero). The declared type of the result of <width bucket function> is the declared type of <width bucket count>.

Access Rules

None.

General Rules

- 1) If the value of one or more <string value expression>s, <datetime value expression>s, <interval value expression>s, and <collection value expression>s that are simply contained in a <numeric value function> is the null value, then the result of the <numeric value function> is the null value.

- 2) If <character position expression> is specified, then let *CVE1* be the value of <character value expression 1> and let *CVE2* be the value of <character value expression 2>.

Case:

- a) If `CHAR_LENGTH(CVE1)` is 0 (zero), then the result is 1 (one).
- b) If <char length units> is specified, then let *CLU* be <char length units>; otherwise, let *CLU* be `CHARACTERS`. If there is at least one value *P* such that

`CVE1 = SUBSTRING (CVE2 FROM P FOR CHAR_LENGTH (CVE1 USING CLU) USING CLU)`

then the result is the least such *P*.

NOTE 133 — The collation used is determined in the normal way.

- c) Otherwise, the result is 0 (zero).

- 3) If <binary position expression> is specified, then

Case:

- a) If the first <binary value expression> has a length of 0 (zero), then the result is 1 (one).
- b) If the value of the first <binary value expression> is equal to an identical-length substring of contiguous octets from the value of the second <binary value expression>, then the result is 1 (one) greater than the number of octets within the value of the second <binary value expression> preceding the start of the first such substring.
- c) Otherwise, the result is 0 (zero).

- 4) If <extract expression> is specified, then

Case:

- a) If <extract field> is a <primary datetime field>, then the result is the value of the datetime field identified by that <primary datetime field> and has the same sign as the <extract source>.

NOTE 134 — If the value of the identified <primary datetime field> is zero or if <extract source> is not an <interval value expression>, then the sign is irrelevant.

- b) Otherwise, let *TZ* be the interval value of the implicit or explicit time zone displacement associated with the <datetime value expression>.

Case:

- i) If <extract field> is `TIMEZONE_HOUR`, then the result is calculated as `EXTRACT (HOUR FROM TZ)`.

- ii) Otherwise, the result is calculated as `EXTRACT (MINUTE FROM TZ)`

- 5) If a <char length expression> is specified, then

Case:

- a) If the character encoding form of <character value expression> is not UTF8, UTF16, or UTF32, then let *S* be the <string value expression>.

Case:

6.27 <numeric value function>

- i) If the most specific type of S is character string, then the result is the number of characters in the value of S .

NOTE 135 — The number of characters in a character string is determined according to the semantics of the character set of that character string.

- ii) Otherwise, the result is OCTET_LENGTH(S).

- b) Otherwise, the result is the number of explicit or implicit <char length units> in <char length expression>, counted in accordance with the definition of those units in the relevant normatively referenced document.
- 6) If an <octet length expression> is specified, then let S be the <string value expression>. Let BL be the number of bits (binary digits) in the value of S . The result of the <octet length expression> is the smallest integer not less than the quotient of the division ($BL/8$).
 - 7) The result of <cardinality expression> is the number of elements of the result of the <collection value expression>.
 - 8) If <absolute value expression> is specified, then let N be the value of the immediately contained <numeric value expression>.

Case:

- a) If N is the null value, then the result is the null value.
 - b) If $N \geq 0$, then the result is N .
 - c) Otherwise, the result is $-1 * N$. If $-1 * N$ is not representable by the result data type, then an exception condition is raised: *data exception — numeric value out of range*.
- 9) If <modulus expression> is specified, then let N be the value of the immediately contained <numeric value expression dividend> and let M be the value of the immediately contained <numeric value expression divisor>.

Case:

- a) If either N or M is the null value, then the result is the null value.
 - b) If M is zero, then an exception condition is raised: *data exception — division by zero*.
 - c) Otherwise, the result is the unique exact numeric value R with scale 0 (zero) such that all of the following are true:
 - i) R has the same sign as N .
 - ii) The absolute value of R is less than the absolute value of M .
 - iii) $N = M * K + R$ for some exact numeric value K with scale 0 (zero).
- 10) If <natural logarithm> is specified, then let V be the value of the simply contained <numeric value expression>.

Case:

- a) If V is the null value, then the result is the null value.
- b) If V is 0 (zero) or negative, then an exception condition is raised: *data exception — invalid argument for natural logarithm*.

c) Otherwise, the result is the natural logarithm of V .

11) If <exponential function> is specified, then let V be the value of the simply contained <numeric value expression>.

Case:

a) If V is the null value, then the result is the null value.

b) Otherwise, the result is e (the base of natural logarithms) raised to the power V . If the result is not representable in the declared type of the result, then an exception condition is raised: *data exception — numeric value out of range*.

12) If <power function> is specified, then let $NVEB$ be the <numeric value expression base>, then let VB be the value of $NVEB$, let $NVEE$ be the <numeric value expression exponent>, and let VE be the value of $NVEE$.

Case:

a) If either VB or VE is the null value, then the result is the null value.

b) If VB is 0 (zero) and VE is negative, then an exception condition is raised: *data exception — invalid argument for power function*.

c) If VB is 0 (zero) and VE is 0 (zero), then the result is 1 (one).

d) If VB is 0 (zero) and VE is positive, then the result is 0 (zero).

e) If VB is negative and VE is not equal to an exact numeric value with scale 0 (zero), then an exception condition is raised: *data exception — invalid argument for power function*.

f) If VB is negative and VE is equal to an exact numeric value with scale 0 (zero) that is an even number, then the result is the result of

$$\text{EXP}(NVEE * \text{LN}(-NVEB))$$

g) If VB is negative and VE is equal to an exact numeric value with scale 0 (zero) that is an odd number, then the result is the result of

$$-\text{EXP}(NVEE * \text{LN}(-NVEB))$$

h) Otherwise, the result is the result of

$$\text{EXP}(NVEE * \text{LN}(NVEB))$$

13) If <floor function> is specified, then let V be the value of the simply contained <numeric value expression> NVE .

Case:

a) If V is the null value, then the result is the null value.

b) Otherwise,

Case:

i) If the most specific type of NVE is exact numeric, then the result is the greatest exact numeric value with scale 0 (zero) that is less than or equal to V . If this result is not representable by the

result data type, then an exception condition is raised: *data exception — numeric value out of range*.

- ii) Otherwise, the result is the greatest whole number that is less than or equal to *V*. If this result is not representable by the result data type, then an exception condition is raised: *data exception — numeric value out of range*.

14) If <ceiling function> is specified, then let *V* be the value of the simply contained <numeric value expression> *NVE*.

Case:

- a) If *V* is the null value, then the result is the null value.
- b) Otherwise,

Case:

- i) If the most specific type of *NVE* is exact numeric, then the result is the least exact numeric value with scale 0 (zero) that is greater than or equal to *V*. If this result is not representable by the result data type, then an exception condition is raised: *data exception — numeric value out of range*.
- ii) Otherwise, the result is the least whole number that is greater than or equal to *V*. If this result is not representable by the result data type, then an exception condition is raised: *data exception — numeric value out of range*.

15) If <width bucket function> is specified, then let *WBO* be the value of <width bucket operand>, let *WBB1* be the value of <width bucket bound 1>, let *WBB2* be the value of <width bucket bound 2>, and let *WBC* be the value of <width bucket count>.

Case:

- a) If any of *WBO*, *WBB1*, *WBB2*, or *WBC* is the null value, then the result is the null value.
- b) If *WBC* is less than or equal to 0 (zero), then an exception condition is raised: *data exception — invalid argument for width bucket function*.
- c) If *WBB1* equals *WBB2*, then an exception condition is raised: *data exception — invalid argument for width bucket function*.
- d) If *WBB1* is less than *WBB2*, then

Case:

- i) If *WBO* is less than *WBB1*, then the result is 0 (zero).
- ii) If *WBO* is greater than or equal to *WBB2*, then the result is *WBC*+1. If the result is not representable in the declared type of the result, then an exception condition is raised: *data exception — numeric value out of range*.
- iii) Otherwise, the result is the greatest exact numeric value with scale 0 (zero) that is less than or equal to $((WBC * (WBO - WBB1) / (WBB2 - WBB1)) + 1)$
- e) If *WBB1* is greater than *WBB2*, then

Case:

- i) If WBO is greater than $WBB1$, then the result is 0 (zero).
- ii) If WBO is less than or equal to $WBB2$, then the result is $WBC+1$. If the result is not representable in the declared type of the result, then an exception condition is raised: *data exception — numeric value out of range*.
- iii) Otherwise, the result is the greatest exact numeric value with scale 0 (zero) that is less than or equal to $((WBC * (WBB1 - WBO) / (WBB1 - WBB2)) + 1)$

16) If <regex occurrences function> is specified, then

- a) Let RSS be the <regex subject string>, let STR be the value of RSS , let PAT be the value of the <XQuery pattern>, let SP be the value of <start position>, let CLU be the <char length units>, and let FL be the value of <XQuery option flag>.
- b) Case:
 - i) If STR , PAT , or FL is the null value, then the result of the <regex occurrences function> is the null value.
 - ii) If SP is less than 1 (one), or greater than the value of

$$\text{CHARACTER_LENGTH} (RSS \text{ USING } CLU)$$
then the result of <regex occurrences function> is -1.
 - iii) If CLU is OCTETS and the SP -th octet of STR is not the first octet of a character, then the result of <regex occurrences function> is implementation-dependent.
 - iv) Otherwise, let $LIST$ be the list of match vectors that is the result of applying the General Rules Subclause 9.16, “XQuery regular expression matching”, with STR as *STRING*, PAT as *PATTERN*, SP as *POSITION*, CLU as *UNITS*, and FL as *FLAG*. The result of <regex occurrences function> is the number of match vectors in $LIST$.

17) If <regex position expression> is specified, then:

- a) Let RSS be the <regex subject string>, let STR be the value of RSS , let PAT be the value of the <XQuery pattern>, let SP be the value of <start position>, let CLU be the <char length units>, let OCC be the value of <regex occurrence>, let CAP be the <regex capture group> and let FL be the value of <XQuery option flag>.
- b) Case:
 - i) If STR , PAT , OCC , CAP , or FL is the null value, then the result of the <regex position expression> is the null value.
 - ii) If OCC is less than 1 (one), then the result of <regex position expression> is 0 (zero).
 - iii) If CAP is less than 0 (zero), or greater than the number of XQuery regular expression parenthesized sub-expressions of PAT , then the result of <regex position expression> is 0 (zero).
 - iv) If SP is less than 1 (one), or greater than the value of

$$\text{CHARACTER_LENGTH} (RSS \text{ USING } CLU)$$
then the result of <regex position expression> is 0 (zero).

- v) If *CLU* is OCTETS and the *SP*-th octet of *STR* is not the first octet of a character, then the result of <regex position expression> is implementation-dependent.
- vi) Otherwise, let *LIST* be the list of match vectors that is the result of applying the General Rules of Subclause 9.16, “XQuery regular expression matching”, with *STR* as *STRING*, *PAT* as *PATTERN*, *SP* as *POSITION*, *CLU* as *UNITS*, and *FL* as *FLAG*.

Case:

- 1) If there are at least *OCC* match vectors in *LIST*, then let *MV* be the *OCC*-th match vector in *LIST*. Let *PL* be *MV*[*CAP*], where *MV*[0] is the first position/length in *MV*. Let *P* be the position of *PL*, and let *L* be the length of *PL*. The result of the <regex position expression> is

Case:

- A) If *P* is 0 (zero), then 0 (zero).
- B) If <regex position start or after> is START, then *P*.
- C) Otherwise, *P* + *L*.

- 2) Otherwise, the result of the <regex position expression> is 0 (zero).

Conformance Rules

- 1) Without Feature S091, “Basic array support”, or Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <cardinality expression>.
- 2) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <extract expression>.
- 3) Feature F411, “Time zone specification”, conforming SQL language shall not contain an <extract expression> that specifies a <time zone field>.
- 4) Without Feature F421, “National character”, conforming SQL language shall not contain a <length expression> that simply contains a <string value expression> that has a declared type of NATIONAL CHARACTER LARGE OBJECT.
- 5) Without Feature T441, “ABS and MOD functions”, conforming language shall not contain an <absolute value expression>.
- 6) Without Feature T441, “ABS and MOD functions”, conforming language shall not contain a <modulus expression>.
- 7) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <natural logarithm>.
- 8) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain an <exponential function>.
- 9) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <power function>.
- 10) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <square root>.

- 11) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <floor function>.
- 12) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <ceiling function>.
- 13) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <width bucket function>.
- 14) Without Feature F842, “OCCURENCES_REGEX function”, conforming SQL language shall not contain <regex occurrences function>.
- 15) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex occurrences function> shall not contain <char length units>.
- 16) Without Feature F843, “POSITION_REGEX function”, conforming SQL language shall not contain <regex position expression>.
- 17) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language , <regex position expression> shall not contain <char length units>.
- 18) Without Feature F847, “Nonconstant regular expression”, in conforming SQL language, <XQuery pattern> and <XQuery option flag> shall be <value specification>s.

6.28 <string value expression>

Function

Specify a character string value or a binary string value.

Format

```

<string value expression> ::=
    <character value expression>
  | <binary value expression>

<character value expression> ::=
    <concatenation>
  | <character factor>

<concatenation> ::=
    <character value expression> <concatenation operator> <character factor>

<character factor> ::=
    <character primary> [ <collate clause> ]

<character primary> ::=
    <value expression primary>
  | <string value function>

<binary value expression> ::=
    <binary concatenation>
  | <binary factor>

<binary factor> ::=
    <binary primary>

<binary primary> ::=
    <value expression primary>
  | <string value function>

<binary concatenation> ::=
    <binary value expression> <concatenation operator> <binary factor>

```

Syntax Rules

- 1) The declared type of a <character primary> shall be character string.
- 2) The declared type of a <character value expression> is the declared type of the simply contained <concatenation> or <character factor>.
- 3) If <concatenation> is specified, then:
 - a) Let *D1* be the declared type of the <character value expression> and let *D2* be the declared type of the <character factor> simply contained in the <concatenation>. Let *D* be the data type determined by applying the Syntax Rules of [Subclause 9.3, “Result of data type combinations”](#), to *D1* and *D2*. Let *CS* be the character set of *D*, let *CO* be the collation of *D*, and let *CD* be the collation derivation of *D*.

- b) Let M be the length in characters of $D1$ plus the length in characters of $D2$. Let VL be the implementation-defined maximum length of variable-length character strings, let LOL be the implementation-defined maximum length of large object character strings, and let FL be the implementation-defined maximum length of fixed-length character strings.

Case:

- i) If the declared type of the <character value expression> or <character factor> is a character large object type, then the declared type of the <concatenation> is a character large object type with character set CS , collation CO , collation derivation CD and maximum length equal to the lesser of M and LOL .
- ii) If the declared type of the <character value expression> or <character factor> is variable-length character string, then the declared type of the <concatenation> is variable-length character string with character set CS , collation CO , collation derivation CD and maximum length equal to the lesser of M and VL .
- iii) If the declared type of the <character value expression> and <character factor> is fixed-length character string, then M shall not be greater than FL and the declared type of the <concatenation> is fixed-length character string with character set CS , collation CO , collation derivation CD and length M .

- 4) If <character factor> is specified, then

Case:

- a) If <collate clause> is specified, then the declared type of the <character factor> is the declared type of the <character primary>, except that the declared type collation is the collation identified by <collate clause>, and its collation derivation is *explicit*.
- b) Otherwise, the declared type of the <character factor> is the declared type of the <character primary>.

- 5) The declared type of <binary primary> shall be binary string.

- 6) Case:

- a) If <binary concatenation> is specified, then let $B1$ be the declared type of the <binary value expression> and let $B2$ be the declared type of the <binary factor>. Let M be the length in octets of $B1$ plus the length in octets of $B2$. Let FL be the implementation-defined maximum length of fixed-length binary strings, let VL be the implementation-defined maximum length of variable-length binary strings, let LOL be the implementation-defined maximum length of binary large object strings.

The declared type of <binary concatenation> is

Case:

- i) If the declared type of the <binary value expression> or <binary factor> is a binary large object type, then binary large object string with maximum length equal to the lesser of M and LOL .
 - ii) If the declared type of the <binary value expression> or <binary factor> is variable-length binary string, then variable-length binary string with maximum length equal to the lesser of M and VL .
 - iii) If the declared type of the <binary value expression> and <binary factor> is fixed-length binary string, then fixed-length binary string with length M , and M shall not be greater than FL .
- b) Otherwise, the declared type of the <binary value expression> is the declared type of the <binary factor>.

Access Rules

None.

General Rules

- 1) If the value of any <character primary> or <binary primary> simply contained in a <string value expression> is the null value, then the result of the <string value expression> is the null value.
- 2) If <concatenation> is specified, then:
 - a) If the character repertoire of <character factor> is UCS, then, in the remainder of this General Rule, the term “length” shall be taken to mean “length in characters”.
 - b) Let *S1* and *S2* be the result of the <character value expression> and <character factor>, respectively.

Case:

- i) If either *S1* or *S2* is the null value, then the result of the <concatenation> is the null value.
- ii) Otherwise:

- 1) Let *S* be the string consisting of *S1* followed by *S2* and let *M* be the length of *S*.
- 2) If the character repertoire of <character factor> is UCS, then *S* is replaced by

Case:

- A) If the <search condition> *S1 IS NORMALIZED AND S2 IS NORMALIZED* evaluates to True, then

NORMALIZE (*S*)

- B) Otherwise, an implementation-defined string.

- 3) Case:

- A) If the most specific type of either *S1* or *S2* is a character large object type, then let *LOL* be the implementation-defined maximum length of large object character strings.

Case:

- I) If *M* is less than or equal to *LOL*, then the result of the <concatenation> is *S* with length *M*.
- II) If *M* is greater than *LOL* and the right-most *M-LOL* characters of *S* are all the <space> character, then the result of the <concatenation> is the first *LOL* characters of *S* with length *LOL*.
- III) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.

- B) If the most specific type of either *S1* or *S2* is variable-length character string, then let *VL* be the implementation-defined maximum length of variable-length character strings.

Case:

- I) If M is less than or equal to VL , then the result of the <concatenation> is S with length M .
 - II) If M is greater than VL and the right-most $M-VL$ characters of S are all the <space> character, then the result of the <concatenation> is the first VL characters of S with length VL .
 - III) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- C) If the most specific types of both $S1$ and $S2$ are fixed-length character string, then the result of the <concatenation> is S .

- 3) If <binary concatenation> is specified, then let $S1$ and $S2$ be the result of the <binary value expression> and <binary factor>, respectively.

Case:

- a) If either $S1$ or $S2$ is the null value, then the result of the <binary concatenation> is the null value.
- b) Otherwise, let S be the string consisting of $S1$ followed by $S2$ and let M be the length in octets of S .

Case:

- i) If the most specific type of either $S1$ or $S2$ is a binary large object type, then let LOL be the implementation-defined maximum length of binary large object strings.

Case:

- 1) If M is less or equal to LOL , then the result of the <binary concatenation> is S with length M .
 - 2) If M is greater than LOL and the right-most $M-LOL$ octets of S are all X'00', then the result of the <binary concatenation> is the first LOL octets of S with length LOL .
 - 3) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- ii) If the most specific type of either $S1$ or $S2$ is variable-length binary string, then let VL be the implementation-defined maximum length of variable-length binary strings.

Case:

- 1) If M is less than or equal to VL , then the result of the <binary concatenation> is S with length M .
 - 2) If M is greater than VL and the right-most $M-VL$ octets of S are all X'00', then the result of the <binary concatenation> is the first VL characters of S with length VL .
 - 3) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- iii) If the most specific types of both $S1$ and $S2$ are fixed-length binary string, then the result of the <binary concatenation> is S .
- 4) If the result of the <character value expression> is a zero-length character string, then it is implementation-defined whether an exception condition is raised: *data exception — zero-length character string*.

Conformance Rules

None.

6.29 <string value function>

*This Subclause is modified by Subclause 6.4, “<string value function>”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 6.8, “<string value function>”, in ISO/IEC 9075-14.*

Function

Specify a function yielding a value of type character string or binary string.

Format

```

09 <string value function> ::=
    <character value function>
  | <binary value function>

14 <character value function> ::=
    <character substring function>
  | <regular expression substring function>
  | <regex substring function>
  | <fold>
  | <transcoding>
  | <character transliteration>
  | <regex transliteration>
  | <trim function>
  | <character overlay function>
  | <normalize function>
  | <specific type method>

<character substring function> ::=
    SUBSTRING <left paren> <character value expression> FROM <start position>
      [ FOR <string length> ] [ USING <char length units> ] <right paren>

<regular expression substring function> ::=
    SUBSTRING <left paren> <character value expression> SIMILAR <character value expression>
      ESCAPE <escape character> <right paren>

<regex substring function> ::=
    SUBSTRING_REGEX <left paren>
      <XQuery pattern> [ FLAG <XQuery option flag> ]
    IN <regex subject string>
      [ FROM <start position> ]
      [ USING <char length units> ]
      [ OCCURRENCE <regex occurrence> ]
      [ GROUP <regex capture group> ]
    <right paren>

<fold> ::=
    { UPPER | LOWER } <left paren> <character value expression> <right paren>

<transcoding> ::=
    CONVERT <left paren> <character value expression>
      USING <transcoding name> <right paren>

<character transliteration> ::=
    TRANSLATE <left paren> <character value expression>

```

6.29 <string value function>

```

    USING <transliteration name> <right paren>

<regex transliteration> ::=
    TRANSLATE_REGEX <left paren>
        <XQuery pattern> [ FLAG <XQuery option flag> ]
        IN <regex subject string>
        [ WITH <XQuery replacement string> ]
        [ FROM <start position> ]
        [ USING <char length units> ]
        [ OCCURRENCE <regex transliteration occurrence> ]
    <left paren>

<XQuery replacement string> ::=
    <character value expression>

<regex transliteration occurrence> ::=
    <regex occurrence>
    | ALL

<trim function> ::=
    TRIM <left paren> <trim operands> <right paren>

<trim operands> ::=
    [ [ <trim specification> ] [ <trim character> ] FROM ] <trim source>

<trim source> ::=
    <character value expression>

<trim specification> ::=
    LEADING
    | TRAILING
    | BOTH

<trim character> ::=
    <character value expression>

<character overlay function> ::=
    OVERLAY <left paren> <character value expression> PLACING <character value expression>
        FROM <start position> [ FOR <string length> ]
        [ USING <char length units> ] <right paren>

<normalize function> ::=
    NORMALIZE <left paren> <character value expression>
        [ <comma> <normal form> [ <comma> <normalize function result length> ] ] <right paren>

<normal form> ::=
    NFC
    | NFD
    | NFKC
    | NFKD

<normalize function result length> ::=
    <character length>
    | <character large object length>

<specific type method> ::=
    <user-defined type value expression> <period> SPECIFICTYPE
        [ <left paren> <right paren> ]

<binary value function> ::=

```

```

    <binary substring function>
  | <binary trim function>
  | <binary overlay function>

<binary substring function> ::=
    SUBSTRING <left paren> <binary value expression> FROM <start position>
        [ FOR <string length> ] <right paren>

<binary trim function> ::=
    TRIM <left paren> <binary trim operands> <right paren>

<binary trim operands> ::=
    [ [ <trim specification> ] [ <trim octet> ] FROM ] <binary trim source>

<binary trim source> ::=
    <binary value expression>

<trim octet> ::=
    <binary value expression>

<binary overlay function> ::=
    OVERLAY <left paren> <binary value expression> PLACING <binary value expression>
        FROM <start position> [ FOR <string length> ] <right paren>

<start position> ::=
    <numeric value expression>

<string length> ::=
    <numeric value expression>

```

Syntax Rules

- 1) 09 The declared type of <string value function> is the declared type of the immediately contained <character value function> or <binary value function>.
- 2) 14 The declared type of <character value function> is the declared type of the immediately contained <character substring function>, <regular expression substring function>, <regex substring function>, <fold>, <transcoding>, <character transliteration>, <regex transliteration>, <trim function>, <character overlay function>, <normalize function>, or <specific type method>.
- 3) The declared type of a <start position>, <string length>, <regex occurrence>, or <regex capture group> shall be exact numeric with scale 0 (zero).
- 4) If <character substring function> *CSF* is specified, then let *DTCVE* be the declared type of the <character value expression> immediately contained in *CSF*. The maximum length, character set, and collation of the declared type *DTCSF* of *CSF* are determined as follows:
 - a) Case:
 - i) If the declared type of <character value expression> is fixed-length character string or variable-length character string, then *DTCSF* is a variable-length character string type with maximum length equal to the fixed length or maximum length of *DTCVE*.
 - ii) Otherwise, the *DTCSF* is a large object character string type with maximum length equal to the maximum length of *DTCVE*.
 - b) The character set and collation of the <character substring function> are those of *DTCVE*.

6.29 <string value function>

- 5) If the character repertoire of <character value expression> is not UCS, then <char length units> shall not be specified.
- 6) If USING <char length units> is not specified, then USING CHARACTERS is implicit.
- 7) If <regular expression substring function> is specified, then:
 - a) The declared types of the <escape character> and the <character value expression>s of the <regular expression substring function> shall be character string with the same character repertoire.
 - b) Case:
 - i) If the declared type of the first <character value expression> is fixed-length character string or variable-length character string, then the declared type of the <regular expression substring function> is variable-length character string with maximum length equal to the maximum variable length of the first <character value expression>.
 - ii) Otherwise, the declared type of the <regular expression substring function> is a character large object type with maximum length equal to the maximum variable length of the first <character value expression>.
 - c) The declared type of the <regular expression substring function> is that of the first <character value expression>.
- 8) If <regex substring function> *RSF* is specified, then:
 - a) If <start position> is not specified, then 1 (one) is implicit.
 - b) Case:
 - i) If <char length units> is specified, then the character repertoire of the <regex subject string> shall be UCS.
 - ii) Otherwise, CHARACTERS is implicit.
 - c) If <regex occurrence> is not specified, then 1 (one) is implicit.
 - d) If <regex capture group> is not specified, then 0 (zero) is implicit.
 - e) If <XQuery option flag> is not specified, then the zero-length string is implicit.
 - f) Let *RSS* be the <regex subject string> immediately contained in *RSF*. Let *DTCVE* be the declared type of *RSS*. The declared type *DTRSF* of *RSF* is determined as follows:
 - i) Case:
 - 1) If *DTCVE* is fixed-length character string or variable-length character string, then *DTRSF* is variable length character string with maximum length equal to the fixed length or maximum length of *DTCVE*.
 - 2) Otherwise, *DTRSF* is a large object character string type with maximum length equal to the maximum length of *DTCVE*.
 - ii) The character set and collation of *DTRSF* are those of *DTCVE*.
- 9) If <fold> is specified, then the declared type of the result of <fold> is that of the <character value expression>.
- 10) If <transcoding> is specified, then:

- a) <transcoding> shall be simply contained in a <value expression> that is immediately contained in a <derived column> that is immediately contained in a <select sublist> or shall immediately contain either a <simple value specification> that is a <host parameter name> or a <value specification> that is a <host parameter specification>.
- b) A <transcoding name> shall identify a transcoding.
- c) Case:
 - i) If the declared type of <character value expression> is fixed-length character string or variable-length character string, then the declared type of the result is variable-length character string with implementation-defined maximum length.
 - ii) Otherwise, the declared type of the result is a character large object type with implementation-defined maximum length.
- d) The character set of the result is an implementation-defined character set *CS* whose character repertoire is the same as the character repertoire of the <character value expression> and whose character encoding form is that determined by the transcoding identified by the <transcoding name>. The declared type collation of the result is the character set collation of *CS*.

11) If <character transliteration> is specified, then:

- a) A <transliteration name> shall identify a character transliteration.
- b) Case:
 - i) If the declared type of <character value expression> is fixed-length character string or variable-length character string, then the declared type of the <character transliteration> is variable-length character string with implementation-defined maximum length.
 - ii) Otherwise, the declared type of the <character transliteration> is a character large object type with implementation-defined maximum length.
- c) The declared type of the <character transliteration> has the character set *CS* that is the target character set of the transliteration. The declared type collation of the result is the character set collation of *CS*.

12) If <regex transliteration> *RT* is specified, then:

- a) If <start position> is not specified, then 1 (one) is implicit.
- b) Case:
 - i) If <char length units> is specified, then the character repertoire of the <regex subject string> shall be UCS.
 - ii) Otherwise, CHARACTERS is implicit.
- c) If <regex transliteration occurrence> is not specified, then ALL is implicit.
- d) If <XQuery replacement string> is not specified, then the zero-length string is implicit.
- e) If <XQuery option flag> is not specified, then the zero-length string is implicit.
- f) Let *RSS* be the <regex subject string> immediately contained in *RT*. Let *DTCVE* be the declared type of *RSS*. The declared type *DTRT* of *RT* is determined as follows:
 - i) Case:

6.29 <string value function>

- 1) If <XQuery replacement string> is not specified, or is a zero-length string, then:
 - A) If *DTCVE* is fixed-length character string or variable-length character string, then *DTRT* is variable length character string with maximum length equal to the fixed length or maximum length of *DTCVE*.
 - B) Otherwise, *DTRT* is a large object character string type with maximum length equal to the maximum length of *DTCVE*.
 - 2) Otherwise:
 - A) If *DTCVE* is fixed-length character string or variable-length character string, then *DTRT* is variable length character string with implementation-defined maximum length.
 - B) Otherwise, *DTRT* is a large object character string type with implementation-defined maximum length.
 - ii) The character set and collation of *DTRT* are those of *DTCVE*.
- 13) If <trim function> is specified, then
- a) Case:
 - i) If FROM is specified, then:
 - 1) Either <trim specification> or <trim character> or both shall be specified.
 - 2) If <trim specification> is not specified, then BOTH is implicit.
 - 3) If <trim character> is not specified, then ' ' is implicit.
 - ii) Otherwise, let *SRC* be <trim source>. TRIM (*SRC*) is equivalent to TRIM (BOTH ' ' FROM *SRC*).
 - b) Case:
 - i) If the declared type of <character value expression> is fixed-length character string or variable-length character string, then the declared type of the <trim function> is variable-length character string with maximum length equal to the fixed length or maximum variable length of the <trim source>.
 - ii) Otherwise, the declared type of the <trim function> is a character large object type with maximum length equal to the maximum variable length of the <trim source>.
 - c) If a <trim character> is specified, then <trim character> and <trim source> shall be comparable.
 - d) The declared type of the <trim function> is that of the <trim source>.
- 14) If <character overlay function> is specified, then:
- a) Let *CV* be the first <character value expression>, let *SP* be the <start position>, and let *RS* be the second <character value expression>.
 - b) If <string length> is specified, then let *SL* be <string length>; otherwise, let *SL* be CHAR_LENGTH(*RS*).
 - c) The <character overlay function> is equivalent to:

SUBSTRING (*CV* FROM 1 FOR *SP* - 1)


```

|| RS
|| SUBSTRING ( CV FROM SP + SL )

```

15) If <normalize function> is specified, then:

a) Case:

- i) If <normal form> is specified, then let *NF* be <normal form>.
- ii) Otherwise, let *NF* be NFC.

b) Case:

- i) If the declared type of <character value expression> is fixed-length character string or variable-length character string, then:
 - 1) <character large object length> shall not be specified.
 - 2) If <normalize function result length> is specified, then let *L* be the value of the <length> contained in <normalize function result length>; otherwise, let *L* be an implementation-defined value of exact numeric type with scale 0 (zero) that is less than or equal to the implementation-defined maximum length of variable-length character strings.
 - 3) The declared type of the <normalize function> is variable-length character string with maximum length equal to *L*, measured in the units of the explicit or implicit <char length units>.
- ii) Otherwise:
 - 1) <character length> shall not be specified.
 - 2) If <normalize function result length> is specified, then let *L* be the value of the <character large object length> contained in <normalize function result length>; otherwise, let *L* be an implementation-defined value of exact numeric type with scale 0 (zero) that is less than or equal to the implementation-defined maximum length of large object character strings.
 - 3) The declared type of the <normalize function> is character large object type with maximum length equal to *L*, measured in the units of the explicit or implicit <char length units>.

16) If <specific type method> is specified, then the declared type of the <specific type method> is variable-length character string with maximum length implementation-defined. The character set of the character string is SQL_IDENTIFIER.

17) 14 The declared type of <binary value function> is the declared type of the immediately contained <binary substring function>, <binary trim function>, or <binary overlay function>.

18) If <binary substring function> *BSF* is specified, then let *DTBVE* be the declared type of the <binary value expression> immediately contained in *BSF*. The declared type *DTBSF* of *BSF* and its maximum length are

Case:

- a) If *DTBVE* is fixed-length binary string or variable-length binary string, then variable-length binary string type with maximum length equal to the fixed or maximum variable length of *DTBVE*.
- b) Otherwise, binary large object string type with maximum length equal to the maximum length of the <binary value expression>.

19) If <binary trim function> is specified, then:

6.29 <string value function>

- a) Case:
 - i) If FROM is specified, then:
 - 1) Either <trim specification> or <trim octet> or both shall be specified.
 - 2) If <trim specification> is not specified, then BOTH is implicit.
 - 3) If <trim octet> is not specified, then X'00' is implicit.
 - ii) Otherwise, let *SRC* be <trim source>. TRIM (*SRC*) is equivalent to TRIM (BOTH X'00' FROM *SRC*).
- b) Case:
 - i) If the declared type of <binary value expression> is fixed-length binary string or variable-length binary string, then the declared type of the <binary trim function> is variable-length binary string type with maximum length equal to the fixed or maximum variable length of the <binary trim source>.
 - ii) Otherwise, the declared type of the <binary trim function> is binary large object string type with maximum length equal to the maximum length of the <binary trim source>.

20) If <binary overlay function> is specified, then:

- a) Let *BV* be the first <binary value expression>, let *SP* be the <start position>, and let *RS* be the second <binary value expression>.
- b) If <string length> is specified, then let *SL* be <string length>; otherwise, let *SL* be OCTET_LENGTH(*RS*).
- c) The <binary overlay function> is equivalent to:

```

SUBSTRING ( BV FROM 1 FOR SP - 1 )
|| RS
|| SUBSTRING ( BV FROM SP + SL )

```

Access Rules

- 1) Case:
 - a) If <string value function> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include USAGE for every transliteration identified by a <transliteration name> contained in the <string value expression>.
 - b) Otherwise, the current privileges shall include USAGE for every transliteration identified by a <transliteration name> contained in the <string value expression>.

General Rules

- 1) The result of <string value function> is the result of the immediately contained <character value function> or <binary value function>.

- 2) 14 The result of <character value function> is the result of the immediately contained <character substring function>, <regular expression substring function>, <regex substring function>, <fold>, <transcoding>, <character transliteration>, <regex transliteration>, <trim function>, <character overlay function>, <normalize function>, or <specific type method>.
- 3) If <character substring function> is specified, then:
 - a) If the character encoding form of <character value expression> is UTF8, UTF16, or UTF32, then, in the remainder of this General Rule, the term “character” shall be taken to mean “unit specified by <char length units>”.
 - b) Let C be the value of the <character value expression>, let LC be the length in characters of C , and let S be the value of the <start position>.
 - c) If <string length> is specified, then let L be the value of <string length> and let E be $S+L$. Otherwise, let E be the larger of $LC + 1$ and S .
 - d) If either C , S , or L is the null value, then the result of the <character substring function> is the null value.
 - e) If E is less than S , then an exception condition is raised: *data exception — substring error*.
 - f) Case:
 - i) If S is greater than LC or if E is less than 1 (one), then the result of the <character substring function> is a zero-length string.
 - ii) Otherwise,
 - 1) Let SI be the larger of S and 1 (one). Let EI be the smaller of E and $LC+1$. Let LI be $EI-SI$.
 - 2) The result of the <character substring function> is a character string containing the LI characters of C starting at character number SI in the same order that the characters appear in C .
- 4) If <regex substring function> is specified, then:
 - a) Let RSS be the <regex subject string>, let STR be the value of RSS , let PAT be the value of the <XQuery pattern>, let SP be the value of <start position>, let CLU be the <char length units>, let OCC be the value of <regex occurrence>, let CAP be the <regex capture group>, and let FL be the value of <XQuery option flag>.
 - b) Case:
 - i) If STR , PAT , OCC , CAP , or FL is the null value, then the result of the <regex substring function> is the null value.
 - ii) If OCC is less than 1 (one), then the result of <regex substring function> is the null value.
 - iii) If CAP is less than 0 (zero), or greater than the number of XQuery regular expression parenthesized sub-expressions of PAT , then the result of <regex substring function> is the null value.
 - iv) If SP is less than 1 (one) or greater than the value of

$$\text{CHARACTER_LENGTH (} RSS \text{ USING } CLU \text{)}$$
then the result of <regex substring function> is the null value.

6.29 <string value function>

- v) If *CLU* is OCTETS and the *SP*-th octet of *STR* is not the first octet of a character, then the result of <regex substring function> is implementation-dependent.
- vi) Let *LIST* be the list of match vectors that is the result of applying the General Rules of Subclause 9.16, “XQuery regular expression matching”, with *STR* as *STRING*, *PAT* as *PATTERN*, *SP* as *POSITION*, *CLU* as *UNITS*, *FL* as *FLAG*.

Case:

- 1) If there are at least *OCC* match vectors in *LIST*, then let *MV* be the *OCC*-th match vector in *LIST*. Let *PL* be *MV*[*CAP*], where *MV*[0] is the first position/length in *MV*. Let *P* be the position of *PL* and let *L* be the length of *PL*. The result of the <regex substring function> is:
 - A) If *P* is 0 (zero), then the null value.
 - B) Otherwise, the substring of *STR* whose position is *P* and whose length is *L*.
 - 2) Otherwise, the result of the <regex substring function> is the null value.
- 5) If <normalize function> is specified, then:
- a) Let *S* be the value of <character value expression>.
 - b) If *S* is the null value, then the result of the <normalize function> is the null value.
 - c) Let *NR* be *S* in the normalized form specified by *NF* in accordance with [UNICODE15].
 - d) Case:
 - i) If the length in characters of *NR* is less than or equal to *L*, then the result of the <normalize function> is *NR*.
 - ii) Otherwise, an exception condition is raised: *data exception — string data, right truncation*.
- 6) If <regular expression substring function> is specified, then:
- a) Let *C* be the result of the first <character value expression>, let *R* be the result of the second <character value expression>, and let *E* be the result of the <escape character>.
 - b) If one or more of *C*, *R*, or *E* is the null value, then the result of the <regular expression substring function> is the null value.
 - c) If the length in characters of *E* is not equal to 1 (one), then an exception condition is raised: *data exception — invalid escape character*.
 - d) If *R* does not contain exactly two occurrences of the two-character sequence consisting of *E*, each immediately followed by <double quote>, then an exception condition is raised: *data exception — invalid use of escape character*.
 - e) Let *R1*, *R2*, and *R3* be the substrings of *R*, such that

$$'R' = 'R1' \ || \ 'E' \ || \ ' "' \ || \ 'R2' \ || \ 'E' \ || \ ' "' \ || \ 'R3'$$
 is True.
 - f) If any one of *R1*, *R2*, or *R3* is not a zero-length string and does not have the format of a <regular expression>, then an exception condition is raised: *data exception — invalid regular expression*.

g) If the predicate

`'C' SIMILAR TO 'R1' || 'R2' || 'R3' ESCAPE 'E'`

is not True, then the result of the <regular expression substring function> is the null value.

h) Otherwise, the result *S* of the <regular expression substring function> is computed as follows:

i) Let *S1* be the shortest initial substring of *C* such that there is a substring *S23* of *C* such that the value of the following <search condition> is True:

`'C' = 'S1' || 'S23' AND
'S1' SIMILAR TO 'R1' ESCAPE 'E' AND
'S23' SIMILAR TO '(R2R3)' ESCAPE 'E'`

ii) Let *S3* be the shortest final substring of *S23* such that there is a substring *S2* of *S23* such that the value of the following <search condition> is True:

`'S23' = 'S2' || 'S3' AND
'S2' SIMILAR TO 'R2' ESCAPE 'E' AND
'S3' SIMILAR TO 'R3' ESCAPE 'E'`

iii) The result of the <regular expression substring function> is *S2*.

7) If <fold> is specified, then:

a) Let *S* be the value of the <character value expression>.

b) If *S* is the null value, then the result of the <fold> is the null value.

c) Let *FRML* be the length or maximum length in characters of the declared type of <fold>.

d) Case:

i) If UPPER is specified, then let *FR* be a copy of *S* in which every lower case character that has a corresponding upper case character or characters in the character set of *S* and every title case character that has a corresponding upper case character or characters in the character set of *S* is replaced by that upper case character or characters.

ii) If LOWER is specified, then let *FR* be a copy of *S* in which every upper case character that has a corresponding lower case character or characters in the character set of *S* and every title case character that has a corresponding lower case character or characters in the character set of *S* is replaced by that lower case character or characters.

e) If the character set of <character factor> is UTF8, UTF16, or UTF32, then *FR* is replaced by

Case:

i) If the <search condition> *S IS NORMALIZED* evaluated to True, then

`NORMALIZE (FR)`

ii) Otherwise, *FR*.

f) Let *FRL* be the length in characters of *FR*.

g) Case:

6.29 <string value function>

- i) If *FRL* is less than or equal to *FRML*, then the result of the <fold> is *FR*. If the declared type of *FR* is fixed-length character string, then the result is padded on the right with $(FRML - FRL)$ <space>s.
- ii) If *FRL* is greater than *FRML*, then the result of the <fold> is the first *FRML* characters of *FR* with length *FRML*. If any of the right-most $(FRL - FRML)$ characters of *FR* are not <space> characters, then a completion condition is raised: *warning — string data, right truncation*.

8) If a <character transliteration> is specified, then

Case:

- a) If the value of <character value expression> is the null value, then the result of the <character transliteration> is the null value.
- b) If <transliteration name> identifies a transliteration descriptor whose indication of how the transliteration is performed specifies an SQL-invoked routine *TR*, then the result of the <character transliteration> is the result of the invocation of *TR* with a single SQL argument that is the <character value expression> contained in the <character transliteration>.
- c) Otherwise, the value of the <character transliteration> is the value returned by the transliteration identified by the <existing transliteration name> specified in the transliteration descriptor of the transliteration identified by <transliteration name>.

9) If <regex transliteration> is specified, then:

- a) Let *RSS* be the <regex subject string>, let *STR* be the value of *RSS*, let *PAT* be the value of the <XQuery pattern>, let *SP* be the value of <start position>, let *CLU* be the <char length units>, let *REP* be the value of <XQuery replacement string>, and let *FL* be the value of <XQuery option flag>.

b) Case:

- i) If *STR*, *PAT*, *REP*, or *FL* is the null value, then the result of the <regex transliteration> is the null value.
- ii) If *SP* is less than 1 (one) or greater than the value of

CHARACTER_LENGTH (*RSS* USING *CLU*)

then the result of <regex transliteration> is the null value.

- iii) If *CLU* is OCTETS and the *SP*-th octet of *STR* is not the first octet of a character, then the result of <regex transliteration> is implementation-dependent.
- iv) Otherwise, let *LIST* be the list of match vectors that is the result of applying the General Rules of Subclause 9.16, “XQuery regular expression matching”, with *STR* as *STRING*, *PAT* as *PATTERN*, *SP* as *POSITION*, *CLU* as *UNITS*, *FL* as *FLAG*. Let *RTO* be the <regex transliteration occurrence>.

Case:

- 1) If *RTO* is <regex occurrence>, then let *OCC* be the value of *RTO*.

Case:

- A) If *OCC* is the null value, then the result of the <regex transliteration> is the null value.

- B) If *OCC* is less than 1 (one) or greater than the number of match vectors in *LIST*, then the result of the <regex transliteration> is *STR*.
 - C) Otherwise, let *MV* be the *OCC*-th match vector in *LIST*. The General Rules of Subclause 9.17, “XQuery regular expression replacement”, are applied with *MV* as *MATCH*, *STR* as *STRING*, *PAT* as *PATTERN*, *REP* as *REPLACEMENT*, and *FL* as *FLAG*, yielding the character string returned from *RESULT*.
- 2) Otherwise, <regex transliteration occurrence> is ALL. Let *N* be the number of match vectors in *LIST*. Let *S_N* be *STR*.
- A) For all *i* between *N* and 1(one) in descending order, let *MV_i* be the *i*-th match vector in *LIST*. The General Rules of Subclause 9.17, “XQuery regular expression replacement”, are applied with *MV_i* as *MATCH*, *S_i* as *STRING*, *PAT* as *PATTERN*, *REP* as *REPLACEMENT*, and *FL* as *FLAG*, yielding *S_{i-1}* returned from *RESULT*.

NOTE 136 — The replacements are applied from the end of *STR* moving to the front to avoid altering the positions of matches within the string. When a replacement is performed, there are no unprocessed replacements later in the string, and consequently the positions of other matches are not disturbed if a substring is replaced by a string of a different length.

- B) The result of <regex transliteration> is *S₀*.

10) If a <transcoding> is specified, then

Case:

- a) If the value of <character value expression> is the null value, then the result of the <transcoding> is the null value.
- b) Otherwise, the value of the <transcoding> is the value of the <character value expression> after the application of the transcoding specified by <transcoding name>.

11) If <trim function> is specified, then:

- a) Let *S* be the value of the <trim source>.
- b) If <trim character> is specified, then let *SC* be the value of <trim character>; otherwise, let *SC* be <space>.
- c) If either *S* or *SC* is the null value, then the result of the <trim function> is the null value.
- d) If the length in characters of *SC* is not 1 (one), then an exception condition is raised: *data exception — trim error*.
- e) Case:
 - i) If BOTH is specified or if no <trim specification> is specified, then the result of the <trim function> is the value of *S* with any leading or trailing characters equal to *SC* removed.
 - ii) If TRAILING is specified, then the result of the <trim function> is the value of *S* with any trailing characters equal to *SC* removed.
 - iii) If LEADING is specified, then the result of the <trim function> is the value of *S* with any leading characters equal to *SC* removed.

12) If <specific type method> is specified, then:

6.29 <string value function>

- a) Let V be the value of the <user-defined type value expression>.
- b) Case:
 - i) If V is the null value, then RV is the null value.
 - ii) Otherwise:
 - 1) Let UDT be the most specific type of V .
 - 2) Let $UDTN$ be the <user-defined type name> of UDT .
 - 3) Let CN be the <catalog name> contained in $UDTN$, let SN be the <unqualified schema name> contained in $UDTN$, and let UN be the <qualified identifier> contained in $UDTN$. Let CND , SND , and UND be CN , SN , and UN , respectively, with every occurrence of <double quote> replaced by <doublequote symbol>. Let RV be:

`"CND" . "SND" . "UND"`

- c) The result of <specific type method> is RV .

13) 14 The result of <binary value function> is the result of the simply contained <binary substring function>, <binary trim function>, or <binary overlay function>.

14) If <binary substring function> is specified, then

- a) Let B be the value of the <binary value expression>, let LB be the length in octets of B , and let S be the value of the <start position>.
- b) If <string length> is specified, then let L be the value of <string length> and let E be $S+L$. Otherwise, let E be the larger of $LB+1$ and S .
- c) If either B , S , or L is the null value, then the result of the <binary substring function> is the null value.
- d) If E is less than S , then an exception condition is raised: *data exception — substring error*.
- e) Case:
 - i) If S is greater than LB or if E is less than 1 (one), then the result of the <binary substring function> is a zero-length string.
 - ii) Otherwise:
 - 1) Let SI be the larger of S and 1 (one). Let EI be the smaller of E and $LB+1$. Let LI be $EI-SI$.
 - 2) The result of the <binary substring function> is a binary string containing LI octets of B starting at octet number SI in the same order that the octets appear in B .

15) If <binary trim function> is specified, then

- a) Let S be the value of the <trim source>.
- b) Let SO be the value of <trim octet>.
- c) If either S or SO is the null value, then the result of the <binary trim function> is the null value.
- d) If the length in octets of SO is not 1 (one), then an exception condition is raised: *data exception — trim error*.

- e) Case:
 - i) If BOTH is specified or if no <trim specification> is specified, then the result of the <binary trim function> is the value of *S* with any leading or trailing octets equal to *SO* removed.
 - ii) If TRAILING is specified, then the result of the <binary trim function> is the value of *S* with any trailing octets equal to *SO* removed.
 - iii) If LEADING is specified, then the result of the <binary trim function> is the value of *S* with any leading octets equal to *SO* removed.
- 16) If the result of <string value expression> is a zero-length character string, then it is implementation-defined whether an exception condition is raised: *data exception — zero-length character string*.

Conformance Rules

- 1) Without Feature T581, “Regular expression substring function”, conforming SQL language shall not contain a <regular expression substring function>.
- 2) Without Feature T312, “OVERLAY function”, conforming SQL language shall not contain a <character overlay function>.
- 3) Without Feature T312, “OVERLAY function”, conforming SQL language shall not contain a <binary overlay function>.
- 4) Without Feature T042, “Extended LOB data type support”, or Feature T022, “Advanced support for BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary value function>.
- 5) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <character transliteration>.
- 6) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transcoding>.
- 7) Without Feature T061, “UCS support”, conforming SQL language shall not contain a <normalize function>.
- 8) Without Feature S261, “Specific type method”, conforming SQL language shall not contain a <specific type method>.
- 9) Without Feature F394, “Optional normal form specification”, conforming SQL language shall not contain <normal form>.
- 10) Without Feature F844, “SUBSTRING_REGEX”, conforming SQL language shall not contain <regex substring function>.
- 11) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex substring function> shall not contain <char length units>.
- 12) Without Feature F845, “TRANSLATE_REGEX”, conforming SQL language shall not contain <regex transliteration>.
- 13) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex transliteration> shall not contain <char length units>.
- 14) Without Feature F847, “Nonconstant regular expressions”, in conforming SQL language, <XQuery pattern>, <XQuery option flag>, and <XQuery replacement string> shall be <value specification>s.

6.30 <datetime value expression>

Function

Specify a datetime value.

Format

```

<datetime value expression> ::=
    <datetime term>
    | <interval value expression> <plus sign> <datetime term>
    | <datetime value expression> <plus sign> <interval term>
    | <datetime value expression> <minus sign> <interval term>

<datetime term> ::=
    <datetime factor>

<datetime factor> ::=
    <datetime primary> [ <time zone> ]

<datetime primary> ::=
    <value expression primary>
    | <datetime value function>

<time zone> ::=
    AT <time zone specifier>

<time zone specifier> ::=
    LOCAL
    | TIME ZONE <interval primary>

```

Syntax Rules

- 1) The declared type of a <datetime primary> shall be datetime.
- 2) If the <datetime value expression> immediately contains neither <plus sign> nor <minus sign>, then the precision of the result of the <datetime value expression> is the precision of the <value expression primary> or <datetime value function> that it simply contains.
- 3) If the declared type of the <datetime primary> is DATE, then <time zone> shall not be specified.
- 4) Case:
 - a) If <time zone> is specified and the declared type of <datetime primary> is TIMESTAMP WITHOUT TIME ZONE or TIME WITHOUT TIME ZONE, then the declared type of <datetime term> is TIMESTAMP WITH TIME ZONE or TIME WITH TIME ZONE, respectively, with the same fractional seconds precision as <datetime primary>.
 - b) Otherwise, the declared type of <datetime term> is the same as the declared type of <datetime primary>.
- 5) If the <datetime value expression> immediately contains either <plus sign> or <minus sign>, then:
 - a) The <interval value expression> or <interval term> shall contain only <primary datetime field>s that are contained within the <datetime value expression> or <datetime term>.

- b) The result of the <datetime value expression> contains the same <primary datetime field>s that are contained in the <datetime value expression> or <datetime term>, with a fractional seconds precision that is the greater of the fractional seconds precisions, if any, of either the <datetime value expression> and <interval term>, or the <datetime term> and <interval value expression> that it simply contains.
- 6) The declared type of the <interval primary> immediately contained in a <time zone specifier> shall be INTERVAL HOUR TO MINUTE.

Access Rules

None.

General Rules

- 1) If the value of any <datetime primary>, <interval value expression>, <datetime value expression>, or <interval term> simply contained in a <datetime value expression> is the null value, then the result of the <datetime value expression> is the null value.
- 2) If <time zone> is specified and the <interval primary> immediately contained in <time zone specifier> is the null value, then the result of the <datetime value expression> is the null value.
- 3) The value of a <datetime primary> is the value of the immediately contained <value expression primary> or <datetime value function>.
- 4) In the following General Rules, arithmetic is performed so as to maintain the integrity of the datetime data type that is the result of the <datetime term> or <datetime value expression>. This may involve carry from or to the immediately next more significant <primary datetime field>. If the data type of the <datetime term> or <datetime value expression> is time with or without time zone, then arithmetic on the HOUR <primary datetime field> is undertaken modulo 24. If the <interval value expression> or <interval term> is a year-month interval, then the DAY field of the result is the same as the DAY field of the <datetime term> or <datetime value expression>.
- 5) The value of a <datetime term> is determined as follows. Let *DT* be the declared type, *DV* the UTC component of the value, and *TZD* the time zone component, if any, of the <datetime primary> simply contained in the <datetime term>, and let *STZD* be the current default time zone displacement of the SQL-session.

Case:

- a) If <time zone> is not specified, then the value of <datetime term> is *DV*.
- b) Otherwise:
 - i) Case:
 - 1) If *DT* is datetime with time zone, then the UTC component of the <datetime term> is *DV*.
 - 2) Otherwise, the UTC component of the <datetime term> is *DV* – *STZD*.
 - ii) Case:
 - 1) If LOCAL is specified, then let *TZ* be *STZD*.
 - 2) If TIME ZONE is specified, then, if the value of the <interval primary> immediately contained in <time zone specifier> is less than INTERVAL - '14:00' or greater than

6.30 <datetime value expression>

INTERVAL + '14:00', then an exception condition is raised: *data exception — invalid time zone displacement value*. Otherwise, let *TZ* be the value of the <interval primary> simply contained in <time zone>.

- iii) The time zone component of the value of the <datetime term> is *TZ*.
- 6) If a <datetime value expression> immediately contains the operator <plus sign> or <minus sign>, then the time zone component, if any, of the result is the same as the time zone component of the immediately contained <datetime term> or <datetime value expression>. The result (if the result type is without time zone) or the UTC component of the result (if the result type has time zone) is effectively evaluated as follows:
 - a) Case:
 - i) If <datetime value expression> immediately contains the operator <plus sign> and the <interval value expression> or <interval term> is not negative, or if <datetime value expression> immediately contains the operator <minus sign> and the <interval term> is negative, then successive <primary datetime field>s of the <interval value expression> or <interval term> are added to the corresponding fields of the <datetime value expression> or <datetime term>.
 - ii) Otherwise, successive <primary datetime field>s of the <interval value expression> or <interval term> are subtracted from the corresponding fields of the <datetime value expression> or <datetime term>.
 - b) If, after the preceding step, any <primary datetime field> of the result is outside the permissible range of values for the field or the result is invalid based on the natural rules for dates and times, then an exception condition is raised: *data exception — datetime field overflow*.

NOTE 137 — For the permissible range of values for <primary datetime field>s, see Table 9, “Valid values for datetime fields”.

Conformance Rules

- 1) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain <datetime value expression> that immediately contains a <plus sign> or a <minus sign>.
- 2) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <time zone>.

6.31 <datetime value function>

Function

Specify a function yielding a value of type datetime.

Format

```
<datetime value function> ::=
    <current date value function>
  | <current time value function>
  | <current timestamp value function>
  | <current local time value function>
  | <current local timestamp value function>

<current date value function> ::=
    CURRENT_DATE

<current time value function> ::=
    CURRENT_TIME [ <left paren> <time precision> <right paren> ]

<current local time value function> ::=
    LOCALTIME [ <left paren> <time precision> <right paren> ]

<current timestamp value function> ::=
    CURRENT_TIMESTAMP [ <left paren> <timestamp precision> <right paren> ]

<current local timestamp value function> ::=
    LOCALTIMESTAMP [ <left paren> <timestamp precision> <right paren> ]
```

Syntax Rules

- 1) The declared type of a <current date value function> is DATE. The declared type of a <current time value function> is TIME WITH TIME ZONE. The declared type of a <current timestamp value function> is TIMESTAMP WITH TIME ZONE.

NOTE 138 — See the Syntax Rules of [Subclause 6.1](#), “<data type>”, for rules governing <time precision> and <timestamp precision>.

- 2) If <time precision> *TP* is specified, then LOCALTIME(*TP*) is equivalent to:

```
CAST (CURRENT_TIME(TP) AS
TIME(TP) WITHOUT TIME
ZONE)
```

Otherwise, LOCALTIME is equivalent to:

```
CAST (CURRENT_TIME AS TIME WITHOUT
TIME ZONE)
```

- 3) If <timestamp precision> *TP* is specified, then LOCALTIMESTAMP(*TP*) is equivalent to:

```
CAST (CURRENT_TIMESTAMP(TP) AS
```

6.31 <datetime value function>

```
TIMESTAMP(TP) WITHOUT TIME
ZONE)
```

Otherwise, LOCALTIMESTAMP is equivalent to:

```
CAST (CURRENT_TIMESTAMP AS TIMESTAMP
WITHOUT TIME ZONE)
```

Access Rules

None.

General Rules

- 1) The <datetime value function>s CURRENT_DATE, CURRENT_TIME, and CURRENT_TIMESTAMP respectively return the current date, current time, and current timestamp; the time and timestamp values are returned with time zone displacement equal to the current default time zone displacement of the SQL-session.
- 2) If specified, <time precision> and <timestamp precision> respectively determine the precision of the time or timestamp value returned.
- 3) Let *S* be an <SQL procedure statement> that is not generally contained in a <triggered action>. All <datetime value function>s that are contained in <value expression>s that are generally contained, without an intervening <routine invocation> whose subject routines do not include an SQL function, either in *S* without an intervening <SQL procedure statement> or in an <SQL procedure statement> contained in the <triggered action> of a trigger activated as a consequence of executing *S*, are effectively evaluated simultaneously. The time of evaluation of a <datetime value function> during the execution of *S* and its activated triggers is implementation-dependent.

NOTE 139 — Activation of triggers is defined in [Subclause 4.38.2, “Trigger execution”](#).

Conformance Rules

- 1) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <current local time value function> that contains a <time precision> that is not 0 (zero).
- 2) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <current local timestamp value function> that contains a <timestamp precision> that is neither 0 (zero) nor 6.
- 3) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <current time value function>.
- 4) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <current timestamp value function>.

6.32 <interval value expression>

Function

Specify an interval value.

Format

```
<interval value expression> ::=
    <interval term>
  | <interval value expression 1> <plus sign> <interval term 1>
  | <interval value expression 1> <minus sign> <interval term 1>
  | <left paren> <datetime value expression> <minus sign> <datetime term> <right paren>
    <interval qualifier>

<interval term> ::=
    <interval factor>
  | <interval term 2> <asterisk> <factor>
  | <interval term 2> <solidus> <factor>
  | <term> <asterisk> <interval factor>

<interval factor> ::=
    [ <sign> ] <interval primary>

<interval primary> ::=
    <value expression primary> [ <interval qualifier> ]
  | <interval value function>

<interval value expression 1> ::=
    <interval value expression>

<interval term 1> ::=
    <interval term>

<interval term 2> ::=
    <interval term>
```

Syntax Rules

- 1) The declared type of an <interval value expression> is interval. The declared type of a <value expression primary> immediately contained in an <interval primary> shall be interval.
- 2) Case:
 - a) If the <interval value expression> simply contains an <interval qualifier> *IQ*, then the declared type of the result is INTERVAL *IQ*.
 - b) If the <interval value expression> is an <interval term>, then the result of the <interval value expression> contains the same interval fields as the <interval primary>. If the <interval primary> contains a seconds field, then the result's fractional seconds precision is the same as the <interval primary>'s fractional seconds precision. The result's <interval leading field precision> is implementation-defined, but shall not be less than the <interval leading field precision> of the <interval primary>.

- c) If <interval term 1> is specified, then the result contains every interval field that is contained in the result of either <interval value expression 1> or <interval term 1>, and, if both contain a seconds field, then the fractional seconds precision of the result is the greater of the two fractional seconds precisions. The <interval leading field precision> is implementation-defined, but shall be sufficient to represent all interval values with the interval fields and <interval leading field precision> of <interval value expression 1> as well as all interval values with the interval fields and <interval leading field precision> of <interval term 1>.

NOTE 140 — Interval fields are effectively defined by Table 4, “Fields in year-month INTERVAL values”, and Table 5, “Fields in day-time INTERVAL values”.

- 3) Case:
 - a) If <interval term 1> is a year-month interval, then <interval value expression 1> shall be a year-month interval.
 - b) If <interval term 1> is a day-time interval, then <interval value expression 1> shall be a day-time interval.
- 4) If <datetime value expression> is specified, then <datetime value expression> and <datetime term> shall be comparable.
- 5) An <interval primary> shall specify <interval qualifier> only if the <interval primary> specifies a <dynamic parameter specification>.

Access Rules

None.

General Rules

- 1) If an <interval term> specifies “<term> * <interval factor>”, then let *T* and *F* be respectively the value of the <term> and the value of the <interval factor>. The result of the <interval term> is the result of $F * T$.
- 2) If the value of any <interval primary>, <datetime value expression>, <datetime term>, or <factor> that is simply contained in an <interval value expression> is the null value, then the result of the <interval value expression> is the null value.
- 3) If *IP* is an <interval primary>, then
 - Case:
 - a) If *IP* immediately contains a <value expression primary> *VEP* and an explicit <interval qualifier> *IQ*, then the value of *IP* is computed by:


```
CAST ( VEP AS INTERVAL IQ )
```
 - b) If *IP* immediately contains a <value expression primary> *VEP*, then the value of *IP* is the value of *VEP*.
 - c) If *IP* is an <interval value function> *IVF*, then the value of *IP* is the value of *IVF*.
- 4) If the <sign> of an <interval factor> is <minus sign>, then the value of the <interval factor> is the negative of the value of the <interval primary>; otherwise, the value of an <interval factor> is the value of the <interval primary>.

5) If <interval term 2> is specified, then:

- a) Let X be the value of <interval term 2> and let Y be the value of <factor>.
- b) Let P and Q be respectively the most significant and least significant <primary datetime field>s of <interval term 2>.
- c) Let E be an exact numeric result of the operation

```
CAST ( CAST ( X AS INTERVAL Q )
AS E1 )
```

where $E1$ is an exact numeric data type of sufficient scale and precision so as to not lose significant digits.

- d) Let OP be the operator $*$ or $/$ specified in the <interval value expression>.
- e) Let I , the result of the <interval value expression> expressed in terms of the <primary datetime field> Q , be the result of

```
CAST ( ( E OP Y ) AS
INTERVAL Q )
```

- f) The result of the <interval value expression> is

```
CAST ( I AS INTERVAL W )
```

where W is an <interval qualifier> identifying the <primary datetime field>s P TO Q , but with <interval leading field precision> such that significant digits are not lost.

6) If <interval term 1> is specified, then let P and Q be respectively the most significant and least significant <primary datetime field>s in <interval term 1> and <interval value expression 1>, let X be the value of <interval value expression 1>, and let Y be the value of <interval term 1>.

- a) Let A be an exact numeric result of the operation

```
CAST ( CAST ( X AS INTERVAL Q )
AS E1 )
```

where $E1$ is an exact numeric data type of sufficient scale and precision so as to not lose significant digits.

- b) Let B be an exact numeric result of the operation

```
CAST ( CAST ( Y AS INTERVAL Q )
AS E2 )
```

where $E2$ is an exact numeric data type of sufficient scale and precision so as to not lose significant digits.

- c) Let OP be the operator $+$ or $-$ specified in the <interval value expression>.
- d) Let I , the result of the <interval value expression> expressed in terms of the <primary datetime field> Q , be the result of:

```
CAST ( ( A OP B ) AS
INTERVAL Q )
```

6.32 <interval value expression>

- e) The result of the <interval value expression> is

```
CAST ( I AS INTERVAL W )
```

where *W* is an <interval qualifier> identifying the <primary datetime field>s *P* TO *Q*, but with <interval leading field precision> such that significant digits are not lost.

- 7) If <datetime value expression> is specified, then let *Y* be the least significant <primary datetime field> specified by <interval qualifier>. Let *DTE* be the <datetime value expression>, let *DT* be the <datetime term>, and let *MSP* be the implementation-defined maximum seconds precision. Evaluation of <interval value expression> proceeds as follows:

- a) Case:

- i) If the declared type of <datetime value expression> is TIME WITH TIME ZONE, then let *A* be the value of:

```
CAST ( DTE AT LOCAL
AS TIME(MSP) WITHOUT
TIME ZONE )
```

- ii) If the declared type of <datetime value expression> is TIMESTAMP WITH TIME ZONE, then let *A* be the value of:

```
CAST ( DTE AT LOCAL
AS TIMESTAMP(MSP) WITHOUT
TIME ZONE )
```

- iii) Otherwise, let *A* be the value of *DTE*.

- b) Case:

- i) If the declared type of <datetime term> is TIME WITH TIME ZONE, then let *B* be the value of:

```
CAST ( DT AT LOCAL
AS TIME(MSP) WITHOUT
TIME ZONE )
```

- ii) If the declared type of <datetime term> is TIMESTAMP WITH TIME ZONE, then let *B* be the value of:

```
CAST ( DT AT LOCAL
AS TIMESTAMP(MSP) WITHOUT
TIME ZONE )
```

- iii) Otherwise, let *B* be the value of *DT*.

- c) *A* and *B* are converted to integer scalars *A2* and *B2* respectively in units *Y* as displacements from some implementation-dependent start datetime.
- d) The result is determined by effectively computing *A2*–*B2* and then converting the difference to an interval using an <interval qualifier> whose <end field> is *Y* and whose <start field> is sufficiently significant to avoid loss of significant digits. The difference of two values of type TIME (with or without time zone) is constrained to be between –24:00:00 and +24:00:00 (excluding each end point); it is implementation-defined which of two non-zero values in this range is the result, although the

computation shall be deterministic. That interval is then converted to an interval using the specified <interval qualifier>, rounding or truncating if necessary. The choice of whether to round or truncate is implementation-defined. If the required number of significant digits exceeds the implementation-defined maximum number of significant digits, then an exception condition is raised: *data exception — interval field overflow*.

Conformance Rules

- 1) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval value expression>.

6.33 <interval value function>

Function

Specify a function yielding a value of type interval.

Format

```
<interval value function> ::=
    <interval absolute value function>

<interval absolute value function> ::=
    ABS <left paren> <interval value expression> <right paren>
```

Syntax Rules

- 1) If <interval absolute value function> is specified, then the declared type of the result is the declared type of the <interval value expression>.

Access Rules

None.

General Rules

- 1) If <interval absolute value function> is specified, then let N be the value of the <interval value expression>.

Case:

 - a) If N is the null value, then the result is the null value.
 - b) If $N \geq 0$ (zero), then the result is N .
 - c) Otherwise, the result is $-1 * N$.

Conformance Rules

- 1) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL shall not contain an <interval value function>.

6.34 <boolean value expression>

Function

Specify a boolean value.

Format

```

<boolean value expression> ::=
    <boolean term>
  | <boolean value expression> OR <boolean term>

<boolean term> ::=
    <boolean factor>
  | <boolean term> AND <boolean factor>

<boolean factor> ::=
    [ NOT ] <boolean test>

<boolean test> ::=
    <boolean primary> [ IS [ NOT ] <truth value> ]

<truth value> ::=
    TRUE
  | FALSE
  | UNKNOWN

<boolean primary> ::=
    <predicate>
  | <boolean predicand>

<boolean predicand> ::=
    <parenthesized boolean value expression>
  | <nonparenthesized value expression primary>

<parenthesized boolean value expression> ::=
    <left paren> <boolean value expression> <right paren>

```

Syntax Rules

- 1) The declared type of a <nonparenthesized value expression primary> shall be boolean.
- 2) If NOT is specified in a <boolean test>, then let *BP* be the contained <boolean primary> and let *TV* be the contained <truth value>. The <boolean test> is equivalent to:

(NOT (*BP* IS *TV*))

- 3) Let *X* denote either a column *C* or the <key word> VALUE. Given a <boolean value expression> *BVE* and *X*, the notion “*BVE* is a readily-known-not-null condition for *X*” is defined as follows.

Case:

- a) If *BVE* is a <predicate> of the form “*RVE* IS NOT NULL”, where *RVE* is a <row value predicand> that is a <row value constructor predicand> that simply contains a <common value expression>,

6.34 <boolean value expression>

<boolean predicand>, or <row value constructor element> that is a <column reference> that references *C*, then *BVE* is a *readily-known-not-null condition* for *C*.

- b) If *BVE* is the <predicate> “VALUE IS NOT NULL”, then *BVE* is a *readily-known-not-null condition* for VALUE.
 - c) Otherwise, *BVE* is not a readily-known-not-null condition for *X*.
- 4) Let *X* denote either a column *C* or the <key word> VALUE. Given a <boolean value expression> *BVE* and *X*, the notion “*BVE* is a *known-not-null condition* for *X*” is defined recursively as follows:
- a) If *BVE* is a <predicate>, then *BVE* is a known-not-null condition for *X* if *BVE* is a readily-known-not-null condition for *X*.
 - b) If the SQL-implementation supports Feature T101, “Enhanced nullability determination”, then *BVE* is also recognized as a known-not-null condition for *X* according to the following recursive rules:
 - i) If *BVE* is a <parenthesized boolean value expression> and the simply contained <boolean value expression> is a known-not-null condition for *X*, then *BVE* is a known-not-null condition for *X*.
 - ii) If *BVE* is a <nonparenthesized value expression primary>, then *BVE* is not a known-not-null condition for *X*.
 - iii) If *BVE* is a <boolean test>, then let *BP* be the <boolean primary> immediately contained in *BVE*. If *BP* is a known-not-null condition for *X*, and <truth value> is not specified, then *BVE* is a known-not-null condition for *X*. Otherwise, *BVE* is not a known-not-null condition for *X*.
 - iv) If *BVE* is of the form “NOT *BT*”, where *BT* is a <boolean test>, then

Case:

 - 1) If *BT* is “*CR* IS NULL”, where *CR* is a column reference that references column *C*, then *BVE* is a known-not-null condition for *C*.
 - 2) If *BT* is “VALUE IS NULL”, then *BVE* is a known-not-null condition for VALUE.
 - 3) Otherwise, *BVE* is not a known-not-null condition for *X*.

NOTE 141 — For simplicity, this rule does not attempt to analyze conditions such as “NOT NOT A IS NULL”, or “NOT (A IS NULL OR NOT (B = 2))”
 - v) If *BVE* is of the form “*BVE1* AND *BVE2*”, then

Case:

 - 1) If either *BVE1* or *BVE2* is a known-not-null condition for *X*, then *BVE* is a known-not-null condition for *X*.
 - 2) Otherwise, *BVE* is not a known-not-null condition for *X*.
 - vi) If *BVE* is of the form “*BVE1* OR *BVE2*”, then *BVE* is not a known-not-null condition for *X*.

NOTE 142 — For simplicity, this rule does not detect cases such as “A IS NOT NULL OR A IS NOT NULL”, which might be classified as a known-not-null condition.
 - c) If *BVE* conforms to an implementation-defined rule that enables the SQL-implementation to correctly infer that, when *BVE* is *True*, then *X* cannot be null, then *BVE* is a *known-not-null condition* for *X*.
- 5) The notion of “retrospectively deterministic” is defined recursively as follows:

- a) A <parenthesized boolean value expression> is retrospectively deterministic if the simply contained <boolean value expression> is retrospectively deterministic.
- b) A <nonparenthesized value expression primary> is retrospectively deterministic if it is not possibly non-deterministic.
- c) A <predicate> *P* is *retrospectively deterministic* if one of the following is true:
 - i) *P* is not possibly non-deterministic.
 - ii) *P* is a <comparison predicate> of the form “ $X < Y$ ”, “ $X \leq Y$ ”, “ $Y > X$ ”, “ $Y \geq X$ ”, “ $X < Y + Z$ ”, “ $X \leq Y + Z$ ”, “ $Y + Z > X$ ”, “ $Y + Z \geq X$ ”, “ $X < Y - Z$ ”, “ $X \leq Y - Z$ ”, “ $Y - Z > X$ ”, or “ $Y - Z \geq X$ ”, where *Y* is CURRENT_DATE, CURRENT_TIMESTAMP or LOCALTIMESTAMP, *X* and *Z* are not possibly non-deterministic <value expression>s, and the declared types of the left and right comparands are either both datetime with time zone or both datetime without time zone.
 - iii) *P* is a <quantified comparison predicate> of the form “ $Y > \text{<quantifier> <table subquery>}$ ”, “ $Y + Z > \text{<quantifier> <table subquery>}$ ”, “ $Y - Z > \text{<quantifier> <table subquery>}$ ”, “ $Y \geq \text{<quantifier> <table subquery>}$ ”, “ $Y + Z \geq \text{<quantifier> <table subquery>}$ ”, or “ $Y - Z \geq \text{<quantifier> <table subquery>}$ ”, where *Y* is CURRENT_DATE, CURRENT_TIMESTAMP or LOCALTIMESTAMP, *Z* is a <value expression> that is not possibly non-deterministic, the <query expression> simply contained in the <table subquery> is not possibly non-deterministic, and the declared types of the left and right comparands are either both datetime with time zone or both datetime without time zone.
 - iv) *P* is a <between predicate> that is transformed into a retrospectively deterministic <boolean value expression>.
- d) A <boolean primary> is retrospectively deterministic if the simply contained <predicate>, <parenthesized boolean value expression> or <nonparenthesized value expression primary> is retrospectively deterministic.
- e) Let *BF* be a <boolean factor>. Let *BP* be the <boolean primary> simply contained in *BF*.
 - i) *BF* is called *negative* if *BF* is of any of the following forms:


```
NOT BP
BP IS FALSE
BP IS NOT TRUE
NOT BP IS NOT FALSE
NOT BP IS TRUE
```
 - ii) *BF* is retrospectively deterministic if one of the following is true:
 - 1) *BF* is negative and *BF* does not generally contain a possibly nondeterministic <value expression>.
 - 2) *BF* is not negative and *BP* is retrospectively deterministic.
- f) A <boolean value expression> is retrospectively deterministic if every simply contained <boolean factor> is retrospectively deterministic.

Access Rules

None.

General Rules

- 1) The result is derived by the application of the specified boolean operators (“AND”, “OR”, “NOT”, and “IS”) to the results derived from each <boolean primary>. If boolean operators are not specified, then the result of the <boolean value expression> is the result of the specified <boolean primary>.
- 2) NOT (*True*) is *False*, NOT (*False*) is *True*, and NOT (*Unknown*) is *Unknown*.
- 3) Table 11, “Truth table for the AND boolean operator”, Table 12, “Truth table for the OR boolean operator”, and Table 13, “Truth table for the IS boolean operator” specify the semantics of AND, OR, and IS, respectively.

Table 11 — Truth table for the AND boolean operator

AND	<i>True</i>	<i>False</i>	<i>Unknown</i>
<i>True</i>	<i>True</i>	<i>False</i>	<i>Unknown</i>
<i>False</i>	<i>False</i>	<i>False</i>	<i>False</i>
<i>Unknown</i>	<i>Unknown</i>	<i>False</i>	<i>Unknown</i>

Table 12 — Truth table for the OR boolean operator

OR	<i>True</i>	<i>False</i>	<i>Unknown</i>
<i>True</i>	<i>True</i>	<i>True</i>	<i>True</i>
<i>False</i>	<i>True</i>	<i>False</i>	<i>Unknown</i>
<i>Unknown</i>	<i>True</i>	<i>Unknown</i>	<i>Unknown</i>

Table 13 — Truth table for the IS boolean operator

IS	TRUE	FALSE	UNKNOWN
<i>True</i>	<i>True</i>	<i>False</i>	<i>False</i>
<i>False</i>	<i>False</i>	<i>True</i>	<i>False</i>
<i>Unknown</i>	<i>False</i>	<i>False</i>	<i>True</i>

Conformance Rules

- 1) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <boolean primary> that simply contains a <nonparenthesized value expression primary>.

- 2) Without Feature F571, “Truth value tests”, conforming SQL language shall not contain a <boolean test> that simply contains a <truth value>.

6.35 <array value expression>

Function

Specify an array value.

Format

```

<array value expression> ::=
    <array concatenation>
  | <array primary>

<array concatenation> ::=
    <array value expression 1> <concatenation operator> <array primary>

<array value expression 1> ::=
    <array value expression>

<array primary> ::=
    <value expression primary>

```

Syntax Rules

- 1) The declared type of the <array value expression> is the declared type of the immediately contained <array concatenation> or <array primary>.
- 2) The declared type of <array primary> is the declared type of the immediately contained <value expression primary>, which shall be an array type.
- 3) If <array concatenation> is specified, then:
 - a) Let *DT* be the data type determined by applying [Subclause 9.3, “Result of data type combinations”](#), to the declared types of <array value expression 1> and <array primary>.
 - b) Let *IMDC* be the implementation-defined maximum cardinality of an array type.
 - c) The declared type of the result of <array concatenation> is an array type whose element type is the element type of *DT* and whose maximum cardinality is the lesser of *IMDC* and the sum of the maximum cardinality of <array value expression 1> and the maximum cardinality of <array primary>.

Access Rules

None.

General Rules

- 1) The value of the result of <array value expression> is the value of the immediately contained <array concatenation> or <array primary>.
- 2) If <array concatenation> is specified, then let *AV1* be the value of <array value expression 1> and let *AV2* be the value of <array primary>.

Case:

- a) If either *AV1* or *AV2* is the null value, then the result of the <array concatenation> is the null value.
- b) If the sum of the cardinality of *AV1* and the cardinality of *AV2* is greater than *IMDC*, then an exception condition is raised: *data exception — array data, right truncation*.
- c) Otherwise, the result is the array comprising every element of *AV1* followed by every element of *AV2*.

Conformance Rules

- 1) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array value expression>.

6.36 <array value constructor>

Function

Specify construction of an array.

Format

```

<array value constructor> ::=
    <array value constructor by enumeration>
  | <array value constructor by query>

<array value constructor by enumeration> ::=
    ARRAY <left bracket or trigraph> <array element list> <right bracket or trigraph>

<array element list> ::=
    <array element> [ { <comma> <array element> }... ]

<array element> ::=
    <value expression>

<array value constructor by query> ::=
    ARRAY <left paren> <query expression> <right paren>

```

Syntax Rules

- 1) The declared type of <array value constructor> is the declared type of the immediately contained <array value constructor by enumeration> or <array value constructor by query>.
- 2) The declared type of the <array value constructor by enumeration> is an array type with element type *DT*, where *DT* is the declared type determined by applying [Subclause 9.3, “Result of data type combinations”](#), to the declared types of the <array element>s immediately contained in the <array element list> of this <array value constructor by enumeration>. The maximum cardinality is the number of <array element>s in the <array element list>, which shall not be greater than the implementation-defined maximum cardinality for array types whose element type is *DT*.
- 3) If <array value constructor by query> is specified, then
 - a) The <query expression> shall be of degree 1 (one). Let *ET* be the declared type of the column in the result of <query expression>.
 - b) The declared type of the <array value constructor by query> is array with element type *ET* and maximum cardinality equal to the implementation-defined maximum cardinality *IMDC* for such array types.

Access Rules

None.

General Rules

- 1) The value of <array value constructor> is the value of the immediately contained <array value constructor by enumeration> or <array value constructor by query>.
- 2) The result of <array value constructor by enumeration> is an array whose i -th element is the value of the i -th <array element> immediately contained in the <array element list>, cast as the data type of DT .
- 3) The result of <array value constructor by query> is determined as follows:
 - a) The <query expression> is evaluated, producing a table T . Let N be the number of rows in T .
 - b) If N is greater than $IMDC$, then an exception condition is raised: *data exception — array data, right truncation*.
 - c) The result of <array value constructor by query> is an array of N elements such that for all i , 1 (one) $\leq i \leq N$, the value of the i -th element is the value of the only column in the i -th row of T .

NOTE 143 — The ordering of the array elements is effectively determined by the General Rules of [Subclause 7.13](#), “<query expression>”.

Conformance Rules

- 1) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array value constructor by enumeration>.
- 2) Without Feature S095, “Array constructors by query”, conforming SQL language shall not contain an <array value constructor by query>.

6.37 <multiset value expression>

Function

Specify a multiset value.

Format

```

<multiset value expression> ::=
    <multiset term>
  | <multiset value expression> MULTISSET UNION [ ALL | DISTINCT ] <multiset term>
  | <multiset value expression> MULTISSET EXCEPT [ ALL | DISTINCT ] <multiset term>

<multiset term> ::=
    <multiset primary>
  | <multiset term> MULTISSET INTERSECT [ ALL | DISTINCT ] <multiset primary>

<multiset primary> ::=
    <multiset value function>
  | <value expression primary>

```

Syntax Rules

- 1) The declared type of a <multiset primary> is the declared type of the immediately contained <multiset value function> or <value expression primary>, which shall be a multiset type.
- 2) If *MI* is a <multiset term> that immediately contains MULTISSET INTERSECT, then let *OP1* be the first operand (the <multiset term>) and let *OP2* be the second operand (the <multiset primary>).
 - a) *OP1* and *OP2* are multiset operands of a multiset element grouping operation. The Syntax Rules of [Subclause 9.11](#), “Multiset element grouping operations”, apply.
 - b) Let *ET1* be the element type of *OP1* and let *ET2* be the element type of *OP2*. Let *ET* be the data type determined by [Subclause 9.3](#), “Result of data type combinations”, using the types *ET1* and *ET2*. The result type of the MULTISSET INTERSECT operation is multiset with element type *ET*.
 - c) If DISTINCT is specified, then let *SQ* be DISTINCT. Otherwise, let *SQ* be ALL.
 - d) *MI* is equivalent to

```

( CASE WHEN OP1 IS NULL OR OP2 IS NULL THEN NULL
  ELSE MULTISSET ( SELECT T1.V
                    FROM UNNEST (OP1) AS T1(V)
                    INTERSECT SQ
                    SELECT T2.V
                    FROM UNNEST (OP2) AS T2(V)
                  )
END )

```

- 3) If *MU* is a <multiset value expression> that immediately contains MULTISSET UNION, then let *OP1* be the first operand (the <multiset value expression>) and let *OP2* be the second operand (the <multiset term>).

- a) If DISTINCT is specified, then *OP1* and *OP2* are multiset operands of a multiset element grouping operation. The Syntax Rules of Subclause 9.11, “Multiset element grouping operations”, apply.
- b) Let *ET1* be the element type of *OP1* and let *ET2* be the element type of *OP2*. Let *ET* be the data type determined by Subclause 9.3, “Result of data type combinations”, using the types *ET1* and *ET2*. The result type of the MULTiset UNION operation is multiset with element type *ET*.
- c) If DISTINCT is specified, then let *SQ* be DISTINCT. Otherwise, let *SQ* be ALL.
- d) *MU* is equivalent to

```
( CASE WHEN OP1 IS NULL OR OP2 IS NULL THEN NULL
  ELSE MULTiset ( SELECT T1.V
                  FROM UNNEST (OP1) AS T1(V)
                UNION SQ
                  SELECT T2.V
                  FROM UNNEST (OP2) AS T2(V)
                )
  END )
```

- 4) If *ME* is a <multiset value expression> that immediately contains MULTiset EXCEPT, then let *OP1* be the first operand (the <multiset term>) and let *OP2* be the second operand (the <multiset primary>).
- a) *OP1* and *OP2* are multiset operands of a multiset element grouping operation. The Syntax Rules of Subclause 9.11, “Multiset element grouping operations”, apply.
- b) Let *ET1* be the element type of *OP1* and let *ET2* be the element type of *OP2*. Let *ET* be the data type determined by Subclause 9.3, “Result of data type combinations”, using the types *ET1* and *ET2*. The result type of the MULTiset EXCEPT operation is multiset with element type *ET*.
- c) If DISTINCT is specified, then let *SQ* be DISTINCT. Otherwise, let *SQ* be ALL.
- d) *ME* is equivalent to

```
( CASE WHEN OP1 IS NULL OR OP2 IS NULL THEN NULL
  ELSE MULTiset ( SELECT T1.V
                  FROM UNNEST (OP1) AS T1(V)
                EXCEPT SQ
                  SELECT T2.V
                  FROM UNNEST (OP2) AS T2(V)
                )
  END )
```

Access Rules

None.

General Rules

- 1) The value of a <multiset primary> is the value of the immediately contained <multiset value function> or <value expression primary>.
- 2) The value of a <multiset term> that is a <multiset primary> is the value of the <multiset primary>.

- 3) The value of a <multiset value expression> that is a <multiset term> is the value of <multiset term>.

Conformance Rules

- 1) Without Feature S275, “Advanced multiset support”, conforming SQL language shall not contain MULTiset UNION, MULTiset INTERSECTION, or MULTiset EXCEPT.

NOTE 144 — If MULTiset UNION DISTINCT, MULTiset INTERSECTION, or MULTiset EXCEPT is specified, then the Conformance Rules of [Subclause 9.11](#), “Multiset element grouping operations”, also apply.

6.38 <multiset value function>

Function

Specify a function yielding a value of a multiset type.

Format

```
<multiset value function> ::=  
    <multiset set function>  
  
<multiset set function> ::=  
    SET <left paren> <multiset value expression> <right paren>
```

Syntax Rules

- 1) Let *MVE* be the <multiset value expression> simply contained in <multiset set function>. *MVE* is a multiset operand of a multiset element grouping operation. The Syntax Rules of [Subclause 9.11](#), “Multiset element grouping operations”, apply.
- 2) The <multiset set function> is equivalent to

```
( CASE WHEN MVE IS NULL THEN NULL  
      ELSE MULTiset ( SELECT DISTINCT M.E  
                      FROM UNNEST (MVE) AS M(E) )  
END )
```

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset value function>.

NOTE 145 — The Conformance Rules of [Subclause 9.11](#), “Multiset element grouping operations”, also apply.

6.39 <multiset value constructor>

Function

Specify construction of a multiset.

Format

```

<multiset value constructor> ::=
    <multiset value constructor by enumeration>
  | <multiset value constructor by query>
  | <table value constructor by query>

<multiset value constructor by enumeration> ::=
    MULTISSET <left bracket or trigraph> <multiset element list> <right bracket or trigraph>

<multiset element list> ::=
    <multiset element> [ { <comma> <multiset element> }... ]

<multiset element> ::=
    <value expression>

<multiset value constructor by query> ::=
    MULTISSET <table subquery>

<table value constructor by query> ::=
    TABLE <table subquery>

```

Syntax Rules

- 1) If <multiset value constructor> immediately contains a <table value constructor by query> *TVCBQ*, then:
 - a) Let *QE* be the <query expression> simply contained in *TVCBQ*.
 - b) *QE* shall not immediately contain an <order by clause>.
 - c) Let *n* be the number of columns in the result of *QE*.
 - d) Let *C*₁, ..., *C*_{*n*} be implementation-dependent identifiers that are all distinct from one another.
 - e) *TVCBQ* is equivalent to

```

MULTISSET ( SELECT ROW ( C1, ..., Cn )
              FROM ( QE ) AS T ( C1, ..., Cn ) )

```

- 2) The declared type of <multiset value constructor> is the declared type of the immediately contained <multiset value constructor by enumeration> or <multiset value constructor by query>.
- 3) The declared type of the <multiset value constructor by enumeration> is a multiset type with element type *DT*, where *DT* is the declared type determined by applying [Subclause 9.3, “Result of data type combinations”](#), to the declared types of the <multiset element>s immediately contained in the <multiset element list> of this <multiset value constructor by enumeration>.
- 4) If <multiset value constructor by query> is specified, then

- a) The <table subquery> *TS* shall be of degree 1 (one).
- b) *TS* shall not immediately contain an <order by clause>.
- c) Let *ET* be the declared type of the column in the result of *TS*.
- d) The declared type of the <multiset value constructor by query> is multiset with element type *ET*.

Access Rules

None.

General Rules

- 1) The value of <multiset value constructor> is the value of the immediately contained <multiset value constructor by enumeration> or <multiset value constructor by query>.
- 2) The result of <multiset value constructor by enumeration> is a multiset whose elements are the values of the <multiset element>s immediately contained in the <multiset element list>, cast as the data type of *DT*.
- 3) If <multiset value constructor by query> is specified, then:
 - a) Let *T* be the value of *TS*. Let *N* be the number of rows in *T*.
 - b) The result of <multiset value constructor by query> is a multiset of *N* elements, with one element for each row of *T*, where the value of each element is the value of the only column in the corresponding row of *T*.

Conformance Rules

- 1) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset value constructor>.
- 2) Without Feature T326, “Table functions”, a <multiset value constructor> shall not contain a <table value constructor by query>.

(Blank page)

7 Query expressions

7.1 <row value constructor>

Function

Specify a value or list of values to be constructed into a row.

Format

```

<row value constructor> ::=
    <common value expression>
  | <boolean value expression>
  | <explicit row value constructor>

<explicit row value constructor> ::=
    <left paren> <row value constructor element> <comma>
      <row value constructor element list> <right paren>
  | ROW <left paren> <row value constructor element list> <right paren>
  | <row subquery>

<row value constructor element list> ::=
    <row value constructor element> [ { <comma> <row value constructor element> }... ]

<row value constructor element> ::=
    <value expression>

<contextually typed row value constructor> ::=
    <common value expression>
  | <boolean value expression>
  | <contextually typed value specification>
  | <left paren> <contextually typed value specification> <right paren>
  | <left paren> <contextually typed row value constructor element> <comma>
      <contextually typed row value constructor element list> <right paren>
  | ROW <left paren> <contextually typed row value constructor element list> <right paren>

<contextually typed row value constructor element list> ::=
    <contextually typed row value constructor element>
      [ { <comma> <contextually typed row value constructor element> }... ]

<contextually typed row value constructor element> ::=
    <value expression>
  | <contextually typed value specification>

<row value constructor predicand> ::=
    <common value expression>
  | <boolean predicand>
  | <explicit row value constructor>

```

Syntax Rules

- 1) If a <row value constructor> is a <common value expression> or a <boolean value expression> *X*, then the <row value constructor> is equivalent to

ROW (*X*)

- 2) If a <row value constructor predicand> is a <common value expression> or a <boolean predicand> *X*, then the <row value constructor predicand> is equivalent to

ROW (*X*)

- 3) Let *ERVC* be an <explicit row value constructor>.

Case:

- a) If *ERVC* simply contains a <row subquery>, then the declared type of *ERVC* is the declared type of that <row subquery>.
 - b) Otherwise, the declared type of *ERVC* is a row type described by a sequence of (<field name>, <data type>) pairs, corresponding in order to each <row value constructor element> *X* simply contained in *ERVC*. The data type is the declared type of *X* and the <field name> is implementation-dependent.
- 4) If a <row value constructor> or <row value constructor predicand> *RVC* is an <explicit row value constructor> *ERVC*, then the declared type of *RVC* is the declared type of *ERVC*.
 - 5) Let *CTRVC* be the <contextually typed row value constructor>.
- a) If *CTRVC* is a <common value expression>, <boolean value expression>, or <contextually typed value specification> *X*, then *CTRVC* is equivalent to:

ROW (*X*)

- b) After the syntactic transformation specified in [SR 5a\)](#) has been performed, if necessary, the declared type of *CTRVC* is a row type described by a sequence of (<field name>, <data type>) pairs, corresponding in order to each <contextually typed row value constructor element> *X* simply contained in *CTRVC*. The <data type> is the declared type of *X* and the <field name> is implementation-dependent.
- 6) The degree of a <row value constructor>, <contextually typed row value constructor>, or <row value constructor predicand> is the degree of its declared type.

Access Rules

None.

General Rules

- 1) The value of a <null specification> is the null value.
- 2) The value of a <default specification> is determined according to the General Rules of [Subclause 11.5](#), “<default clause>”.
- 3) The value of an <empty specification> is an empty collection.

4) Case:

- a) If a <row value constructor>, <row value constructor predicand>, or <contextually typed row value constructor> immediately contains a <common value expression>, <boolean value expression>, or <contextually typed row value constructor element> X, then the result of the <row value constructor>, <row value constructor predicand>, or <contextually typed row value constructor> is a row containing a single column whose value is the value of X.
- b) If an <explicit row value constructor> is specified, then the result of the <row value constructor> or <row value constructor predicand> is a row of columns, the value of whose *i*-th column is the value of the *i*-th <row value constructor element> simply contained in the <explicit row value constructor>.
- c) If a <contextually typed row value constructor element list> is specified, then the result of the <contextually typed row value constructor> is a row of columns, the value of whose *i*-th column is the value of the *i*-th <contextually typed row value constructor element> in the <contextually typed row value constructor element list>.

Conformance Rules

- 1) Without Feature T051, “Row types”, conforming SQL language shall not contain an <explicit row value constructor> that immediately contains ROW.
- 2) Without Feature T051, “Row types”, conforming SQL language shall not contain a <contextually typed row value constructor> that immediately contains ROW.
- 3) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain an <explicit row value constructor> that is not simply contained in a <table value constructor> and that contains more than one <row value constructor element>.
- 4) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain an <explicit row value constructor> that is a <row subquery>.
- 5) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <row value constructor predicand> that immediately contains a <boolean predicand>.
- 6) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain a <contextually typed row value constructor> that is not simply contained in a <contextually typed table value constructor> and that contains more than one <row value constructor element>.
- 7) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain a <contextually typed row value constructor> that is a <row subquery>.

7.2 <row value expression>

Function

Specify a row value.

Format

```
<row value expression> ::=  
    <row value special case>  
    | <explicit row value constructor>  
  
<table row value expression> ::=  
    <row value special case>  
    | <row value constructor>  
  
<contextually typed row value expression> ::=  
    <row value special case>  
    | <contextually typed row value constructor>  
  
<row value predicand> ::=  
    <row value special case>  
    | <row value constructor predicand>  
  
<row value special case> ::=  
    <nonparenthesized value expression primary>
```

Syntax Rules

- 1) The declared type of a <row value special case> shall be a row type.
- 2) The declared type of a <row value expression> is the declared type of the immediately contained <row value special case> or <explicit row value constructor>.
- 3) The declared type of a <table row value expression> is the declared type of the immediately contained <row value special case> or <row value constructor>.
- 4) The declared type of a <contextually typed row value expression> is the declared type of the immediately contained <row value special case> or <contextually typed row value constructor>. The declared type of a <row value predicand> is the declared type of the immediately contained <row value special case> or <row value constructor predicand>.

Access Rules

None.

General Rules

- 1) A <row value special case> specifies the row value denoted by the <nonparenthesized value expression primary>.

- 2) A <row value expression> specifies the row value denoted by the <row value special case> or <explicit row value constructor>.
- 3) A <table row value expression> specifies the row value denoted by the <row value special case> or <row value constructor>.
- 4) A <contextually typed row value expression> specifies the row value denoted by the <row value special case> or <contextually typed row value constructor>.
- 5) A <row value predicand> specifies the row value denoted by the <row value special case> or <row value constructor predicand>.

Conformance Rules

- 1) Without Feature T051, “Row types”, conforming SQL language shall not contain a <row value special case>.

7.3 <table value constructor>

Function

Specify a set of <row value expression>s to be constructed into a table.

Format

```

<table value constructor> ::=
    VALUES <row value expression list>

<row value expression list> ::=
    <table row value expression> [ { <comma> <table row value expression> }... ]

<contextually typed table value constructor> ::=
    VALUES <contextually typed row value expression list>

<contextually typed row value expression list> ::=
    <contextually typed row value expression>
    [ { <comma> <contextually typed row value expression> }... ]

```

Syntax Rules

- 1) All <table row value expression>s immediately contained in a <row value expression list> shall be of the same degree.
- 2) All <contextually typed row value expression>s immediately contained in a <contextually typed row value expression list> shall be of the same degree.
- 3) A <table value constructor> or a <contextually typed table value constructor> is *possibly non-deterministic* if it generally contains a possibly non-deterministic <value expression>.
- 4) Let *TVC* be some <table value constructor> consisting of n <table row value expression>s or some <contextually typed table value constructor> consisting of n <contextually typed row value expression>s. Let RVE_i , $1 \text{ (one)} \leq i \leq n$, denote the i -th <table row value expression> or the i -th <contextually typed row value expression>. The row type of *TVC* is determined by applying [Subclause 9.3, “Result of data type combinations”](#), to the row types RVE_i , $1 \text{ (one)} \leq i \leq n$. The column names are implementation-dependent.

Access Rules

None.

General Rules

- 1) If the result of any <table row value expression> or <contextually typed row value expression> is the null value, then an exception condition is raised: *data exception — null row not permitted in table*.
- 2) The result T of a <table value constructor> or <contextually typed table value constructor> *TVC* is a table whose cardinality is the number of <table row value expression>s or the number of <contextually typed row value expression>s in *TVC*. If R is the result of n such expressions, then R occurs n times in T .

Conformance Rules

- 1) Without Feature F641, “Row and table constructors”, in conforming SQL language, the <contextually typed row value expression list> of a <contextually typed table value constructor> shall contain exactly one <contextually typed row value constructor> *RVE*. *RVE* shall be of the form “(<contextually typed row value constructor element list>)”.
- 2) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain a <table value constructor>.

7.4 <table expression>

Function

Specify a table or a grouped table.

Format

```
<table expression> ::=
  <from clause>
    [ <where clause> ]
    [ <group by clause> ]
    [ <having clause> ]
    [ <window clause> ]
```

Syntax Rules

- 1) The result of a <table expression> is a derived table whose row type is the row type of the result of the immediately contained <from clause>, together with the window structure descriptors defined by the <window clause>, if specified.
- 2) Let *C* be some column. Let *TE* be the <table expression>. *C* is an underlying column of *TE* if and only if *C* is an underlying column of some column reference contained in *TE*.

Access Rules

None.

General Rules

- 1) If all optional clauses are omitted, then the result of the <table expression> is the same as the result of the <from clause>. Otherwise, each specified clause is applied to the result of the previously specified clause and the result of the <table expression> is the result of the application of the last specified clause.

Conformance Rules

None.

7.5 <from clause>

Function

Specify a table derived from one or more tables.

Format

```
<from clause> ::=
  FROM <table reference list>

<table reference list> ::=
  <table reference> [ { <comma> <table reference> }... ]
```

Syntax Rules

- 1) Let *TRL* be the ordering of <table reference list>. No element TR_i in *TRL* shall contain an outer reference to an element TR_j , where $i \leq j$.
- 2) Case:
 - a) If the <table reference list> immediately contains a single <table reference>, then the descriptor of the result of the <table reference list> is the same as the descriptor of the table identified by that <table reference>. The row type *RT* of the result of the <table reference list> is the row type of the table identified by the <table reference>.
 - b) If the <table reference list> immediately contains more than one <table reference>, then the descriptors of the columns of the result of the <table reference list> are the descriptors of the columns of the tables identified by the <table reference>s, in the order in which the <table reference>s appear in the <table reference list> and in the order in which the columns are defined within each table. The row type *RT* of the result of the <table reference list> is determined by the sequence *SCD* of column descriptors of the result as follows:
 - i) Let n be the number of column descriptors in *SCD*. *RT* has n fields.
 - ii) For i ranging from 1 (one) to n , the field name of the i -th field descriptor in *RT* is the column name included in the i -th column descriptor in *SCD*.
 - iii) For i ranging from 1 (one) to n , the data type descriptor of the i -th field descriptor in *RT* is

Case:

 - 1) If the i -th column descriptor in *SCD* includes a domain name *DN*, then the data type descriptor included in the descriptor of the domain identified by *DN*.
 - 2) Otherwise, the data type descriptor included in the i -th column descriptor in *SCD*.
- 3) The descriptor of the result of the <from clause> is the same as the descriptor of the result of the <table reference list>.

Access Rules

None.

General Rules

- 1) Let *TRLR* be the result of *TRL*.

Case:

- a) If *TRL* simply contains a single <table reference> *TR*, then *TRLR* is the result of *TR*.
- b) If *TRL* simply contains n <table reference>s, where $n > 1$, then let *TRL-P* be the <table reference list> formed by taking the first $n-1$ elements of *TRL* in order, let *TRL-L* be the last element of *TRL*, and let *TRLR-P* be the result of *TRL-P*. If *TRLR-P* contains m rows, $m \geq 1$ (one), then for every row R_i , 1 (one) $\leq i \leq m$, let *TRLR-L_i* be the corresponding evaluation of *TRL-L* under all outer references contained in *TRL-L*. Let *SUBR_i* be the table containing every row formed by concatenating R_i with some row of *TRLR-L_i*. Every row *RR* in *SUBR_i* is a row in *TRLR*, and the number of occurrences of *RR* in *TRLR* is the sum of the numbers of occurrences of *RR* in every occurrence of *SUBR_i*.

The result of the <table reference list> is *TRLR* with the columns reordered according to the ordering of the descriptors of the columns of the <table reference list>.

- 2) The result of the <from clause> is *TRLR*.

Conformance Rules

None.

7.6 <table reference>

*This Subclause is modified by Subclause 7.1, “<table reference>”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 7.1, “<table reference>”, in ISO/IEC 9075-14.*

Function

Reference a table.

Format

```
<table reference> ::=
    <table factor>
  | <joined table>

<table factor> ::=
    <table primary> [ <sample clause> ]

<sample clause> ::=
    TABLESAMPLE <sample method> <left paren> <sample percentage> <right paren>
    [ <repeatable clause> ]

<sample method> ::=
    BERNOULLI
  | SYSTEM

<repeatable clause> ::=
    REPEATABLE <left paren> <repeat argument> <right paren>

<sample percentage> ::=
    <numeric value expression>

<repeat argument> ::=
    <numeric value expression>

14 <table primary> ::=
    <table or query name> [ [ AS ] <correlation name>
    [ <left paren> <derived column list> <right paren> ] ]
  | <derived table> [ AS ] <correlation name>
    [ <left paren> <derived column list> <right paren> ]
  | <lateral derived table> [ AS ] <correlation name>
    [ <left paren> <derived column list> <right paren> ]
  | <collection derived table> [ AS ] <correlation name>
    [ <left paren> <derived column list> <right paren> ]
  | <table function derived table> [ AS ] <correlation name>
    [ <left paren> <derived column list> <right paren> ]
  | <only spec> [ [ AS ] <correlation name>
    [ <left paren> <derived column list> <right paren> ] ]
  | <parenthesized joined table>

<only spec> ::=
    ONLY <left paren> <table or query name> <right paren>

<lateral derived table> ::=
    LATERAL <table subquery>
```

7.6 <table reference>

```

<collection derived table> ::=
    UNNEST <left paren> <collection value expression>
        [ { <comma> <collection value expression> }... ] <right paren>
        [ WITH ORDINALITY ]

<table function derived table> ::=
    TABLE <left paren> <collection value expression> <right paren>

<derived table> ::=
    <table subquery>

<table or query name> ::=
    <table name>
    | <transition table name>
    | <query name>

<derived column list> ::=
    <column name list>

<column name list> ::=
    <column name> [ { <comma> <column name> }... ]

<parenthesized joined table> ::=
    <left paren> <parenthesized joined table> <right paren>
    | <left paren> <joined table> <right paren>

```

Syntax Rules

- 1) The declared type of <repeat argument> shall be an exact numeric type with scale 0 (zero).
- 2) If a <table primary> *TP* simply contains a <table function derived table> *TFDT*, then:
 - a) The <collection value expression> immediately contained in *TFDT* shall be a <routine invocation>.
 - b) Let *CN* be the <correlation name> simply contained in *TP*.
 - c) Let *CVE* be the <collection value expression> simply contained in *TP*.
 - d) Case:
 - i) If *TP* specifies a <derived column list> *DCL*, then let *TFDCL* be


```
( DCL )
```
 - ii) Otherwise, let *TFDCL* be a zero-length string.
 - e) *TP* is equivalent to the <table primary>


```
UNNEST ( CVE ) AS CN TFDCL
```
- 3) If a <table primary> *TP* simply contains a <collection derived table> *CDT*, then:
 - a) Let *NCV* be the number of <collection value expression>s simply contained in *CDT*.
 - b) Let *CVE_j*, $1 \leq j \leq NCV$, be an enumeration of the <collection value expression>s simply contained in *CDT*, in order from left to right.

- c) Let ET_j , $1 \text{ (one)} \leq j \leq NCV$, be the element type of the declared type of CVE_j .
- d) Let CN be the <correlation name> simply contained in TP , and let $TEMP$ be an <identifier> that is not equivalent to CN nor to any other <identifier> contained in TP .

e) Case:

- i) If the declared type of any CVE_j , $1 \text{ (one)} \leq j \leq NCV$, is a multiset, then NCV shall be 1 (one) and WITH ORDINALITY shall not be specified. Let $IMDC_1$ be the implementation-defined maximum cardinality of an array whose declared element type is ET_1 . Let C_1 be

```
( CAST ( CVE1 AS ET1 ARRAY[IMDC1] ) )
```

- ii) Otherwise, let C_j , $1 \text{ (one)} \leq j \leq NCV$, be CVE_j .

f) Let $CARD_j$, $1 \text{ (one)} \leq j \leq NCV$, be

```
COALESCE( CARDINALITY( Cj ), 0 )
```

g) Let $MCARD_1$ be $CARD_1$. Let $MCARD_j$, $2 \leq j \leq NCV$, be

```
CASE
  WHEN CARDj > CARDj-1
    THEN CARDj
  ELSE CARDj-1
END
```

h) Let $MAXCARD$ be $MCARD_{NCV}$.

i) Let N_j , $1 \text{ (one)} \leq j \leq NCV$, and $NORD$ be $NCV+1$ <column name>s that are not equivalent to one another nor to CN , $TEMP$, or any other <identifier> contained in TP .

j) Let $ELTl_j$, $1 \text{ (one)} \leq j \leq NCV$, be

```
CASE
  WHEN 1 <= CARDINALITY( Cj )
    THEN Cj[1]
  ELSE NULL
END
```

k) Let $ELTNORD_j$, $1 \text{ (one)} \leq j \leq NCV$, be

```
CASE
  WHEN NORD < CARDINALITY( Cj )
    THEN Cj[NORD+1]
  ELSE NULL
END
```

l) Let $RECQP$ be:

```
WITH RECURSIVE TEMP(N1, ..., NNCV, NORD) AS
( SELECT ELTl1 AS N1, ..., ELTlNCV AS NNCV, 1 AS NORD
```

```

        FROM (VALUES(1)) AS CN
        WHERE 0 < MAXCARD
    UNION
        SELECT ELTNORD1 AS N1, . . . , ELTNORDNCV AS NNCV, NORD+1 AS NORD
        FROM TEMP
        WHERE NORD < MAXCARD
    )

```

m) Case:

i) If *TP* specifies a <derived column list> *DCL*, then:

1) For all *j* between 1 (one) and *NCV*,

Case:

A) If *ET_j* is a row type, then let *DET_j* be the degree of *ET_j*.

B) Otherwise, let *DET_j* be 1 (one).

2) Case:

A) If *CDT* specifies WITH ORDINALITY, then *DCL* shall contain *DET₁* + ... + *DET_{NCV}* + 1 <column name>s.

B) Otherwise, *DCL* shall contain *DET₁* + ... + *DET_{NCV}* <column name>s.

3) Let *PDCLP* be

(*DCL*)

ii) Otherwise,

Case:

1) If any *ET_j*, 1 (one) ≤ *j* ≤ *NCV*, is a row type, then:

A) For each *ET_j* that is a row type:

I) Let *DET(j)* be the degree of *ET_j*.

II) Let *FN_{j,i}*, 1 (one) ≤ *i* ≤ *DET(j)*, be the name of the *i*-th field in *ET_j*.

B) For each *ET_j* that is not a row type, let *DET(j)* be 1 (one) and let *FN_{j,1}* be an <identifier> that is not equivalent to any other <identifier> contained in *TP*.

C) Case:

I) If *CDT* specifies WITH ORDINALITY, then let *PDCLP* be:

```

    ( FN1,1, FN1,2, . . . , FN1,DET(1),
      . . . ,
      FNNCV,1, FNNCV,2, . . . , FNNCV,DET(NCV), NORD
    )

```

II) Otherwise, let *PDCLP* be:

$$\begin{aligned}
 & (FN_{1,1}, FN_{1,2}, \dots, FN_{1,DET(1)}, \\
 & \dots, \\
 & FN_{NCV,1}, FN_{NCV,2}, \dots, FN_{NCV,DET(NCV)} \\
 &)
 \end{aligned}$$

2) Otherwise, let *PDCLP* be a zero-length string.

n) If *ET_j*, 1 (one) ≤ *j* ≤ *NCV*, is a row type, then let *SLE_j* be (*N_j*) . *; otherwise, let *SLE_j* be *N_j*.

Case:

i) If *CDT* specifies WITH ORDINALITY, then let *ELDT* be:

```

LATERAL ( RECQP SELECT SLE1, ..., SLENCV, NORD
          FROM TEMP ) AS CN PDCLP

```

ii) Otherwise, let *ELDT* be:

```

LATERAL ( RECQP SELECT SLE1, ..., SLENCV
          FROM TEMP ) AS CN PDCLP

```

o) *TP* is equivalent to the <table primary> *ELDT*.

4) Let *TR* be a <table reference>.

Case:

a) If *TR* simply contains a <joined table>, then let *JT* be that <joined table>.

b) Otherwise, let *TF* be the <table factor> that is immediately contained in *TR*, and let *TP* be the <table primary> that is immediately contained in *TF*.

5) If *TF* simply contains a <correlation name>, then let *RV* be that <correlation name>; otherwise, let *RV* be the <table or query name> simply contained in *TF*. *RV* is a range variable. *RV* is *exposed* by *TF* and by *TR*.

NOTE 146 — “range variable” is defined in Subclause 4.14.6, “Operations involving tables”.

NOTE 147 — Thus a range variable is a BNF nonterminal (either a <correlation name> or a <table or query name>). It is possible for an <SQL procedure statement> to contain several range variables whose <identifier>s are equivalent, which must nevertheless be regarded as distinct range variables.

6) Case:

a) If *TR* is contained in a <from clause> *FC* with no intervening <query expression>, then the *scope clause SC* of *TR* is the <select statement: single row> or innermost <query specification> that contains *FC*. The scope of a range variable of *TR* is the <select list>, <where clause>, <group by clause>, <having clause>, and <window clause> of *SC*, together with every <lateral derived table> that is simply contained in *FC* and is preceded by *TR*, and every <collection derived table> that is simply contained in *FC* and is preceded by *TR*, and the <join condition> of all <joined table>s contained in *SC* that contain *TR*. If *SC* is the <query specification> that is the <query expression body> of a simple table query *STQ*, then the scope of a range variable of *TR* also includes the <order by clause> of *STQ*.

NOTE 148 — “simple table query” is defined in Subclause 7.13, “<query expression>”.

b) If *TF* is simply contained in a <merge statement> *MS*, then the *scope clause SC* of *TR* is *MS*. The scope of the range variable of *TF* and of *TR* is the <search condition>, <set clause list>, and <merge insert value list> of *SC*.

NOTE 149 — Subclause 14.12, “<merge statement>”, does not allow *TR* to directly contain a <joined table>.

- c) Otherwise, the *scope clause* *SC* of *TR* is the outermost <joined table> that contains *TR* with no intervening <query expression>. The scope of any range variable of *TR* is the <join condition> of *SC* and of all <joined table>s contained in *SC* that contain *TR*.
- 7) Let *RV* be a range variable that is exposed by *TR*. Let *RV1* be a range variable that is exposed by a <table reference> *TR1* that has the same scope clause as *TR*.

Case:

- a) If *RV* is a <table name>, then

Case:

- i) If *RV1* is a <table name>, then *RV1* shall not be equivalent to *RV*.
- ii) Otherwise, *RV1* shall not be equivalent to the <qualified identifier> of *RV*.

- b) Otherwise,

Case:

- i) If *RV1* is a <table name>, then the <qualified identifier> of *RV1* shall not be equivalent to *RV*.
- ii) Otherwise, *RV1* shall not be equivalent to *RV*.

- 8) A <table or query name> simply contained in *TF* has a scope clause and scope defined by *TF* if and only if the <table or query name> is exposed by *TF* and by *TR*.
- 9) If *TP* simply contains <table or query name> *TOQN*, then

Case:

- a) If *TOQN* is an <identifier> that is equivalent to a <query name> *QN*, then let *WLE* be the <with list element> simply contained in the <query expression> that simply contains *TP* such that the <query name> *QNI* simply contained in *WLE* is equivalent to *QN* and *QNI* is the innermost query name in scope. Let the *table specified by the <query name>* be the result of *WLE*.

NOTE 150 — “query name in scope” is defined in Subclause 7.13, “<query expression>”.

- b) If *TOQN* is an <identifier> that is equivalent to a <transition table name> that is in scope, then let the *table specified by the <transition table name>* be the table identified by *TOQN*.

NOTE 151 — The scope of a <transition table name> is defined in Subclause 11.40, “<trigger definition>”.

- c) Otherwise, let the *table specified by the <table name>* be the table identified by the <table name> simply contained in *TP*.

NOTE 152 — The preceding cases disambiguate whether *TOQN* is interpreted as a <query name>, <transition table name>, or <table name>.

- 10) If *TP* simply contains <only spec> *OS* and the table identified by the <table or query name> *TN* is not a typed table, then *OS* is equivalent to *TN*.
- 11) No <column name> shall be specified more than once in a <derived column list>.
- 12) If a <derived column list> is specified in *TP*, then the number of <column name>s in the <derived column list> shall be the same as the degree of the table specified by the <derived table>, the <lateral derived table>, or the <table or query name> simply contained in *TP*, and the name of the *i*-th column of that

<derived table> or <lateral derived table> or the effective name of the *i*-th column of that <table or query name> is the *i*-th <column name> in that <derived column list>.

13) The row type of a <lateral derived table> is the row type of the simply contained <query expression>.

14) Case:

- a) If no <derived column list> is specified in *TR*, then the row type *RT* of *TR* is the row type of its simply contained <table or query name>, <derived table>, <lateral derived table>, or <joined table>.
- b) Otherwise, the row type *RT* of *TR* is described by a sequence of (<field name>, <data type>) pairs, where the <field name> in the *i*-th pair is the *i*-th <column name> in the <derived column list> and the <data type> in the *i*-th pair is the declared type of the *i*-th column of the <derived table>, <joined table>, <lateral derived table>, or of the table identified by the <table or query name> simply contained in *TR*.

15) A <derived table> or <lateral derived table> is an *updatable derived table* if and only if the <query expression> simply contained in the <derived table> or <lateral derived table> is updatable.

16) A <derived table> or <lateral derived table> is a *simply updatable derived table* if and only if the <query expression> simply contained in the <derived table> or <lateral derived table> simply is updatable.

17) A <derived table> or <lateral derived table> is an *insertable-into derived table* if and only if the <query expression> simply contained in the <derived table> or <lateral derived table> is insertable-into.

18) A <collection derived table> is not updatable and is not simply updatable.

19) If a <table reference> *TR* immediately contains a <table factor> *TF*, then

Case:

- a) If *TF* simply contains a <table name> that identifies a base table, then every column of the table identified by *TF* is called an *updatable column* of *TR*.
- b) If *TF* simply contains a <table name> that identifies a view, then every updatable column of the view identified by *TF* is called an *updatable column* of *TR*.
- c) If *TF* simply contains a <derived table> or <lateral derived table>, then every updatable column of the table identified by the <query expression> simply contained in <derived table> or <lateral derived table> is called an *updatable column* of *TR*.

20) If a <table reference> *TR* immediately contains a <table factor> and the <table or query name> simply contained in *TR* immediately contains a <table name> *TN*, then let *T* be the table identified by *TN*. The schema identified by the explicit or implicit qualifier of *TN* shall include the descriptor of *T*.

21) A <table name> is *possibly non-deterministic* if the table identified by the <table name> is a viewed table, and the original <query expression> in the view descriptor identified by the <table name> is possibly non-deterministic.

22) A <query name> is *possibly non-deterministic* if the <query expression> identified by the <query name> is possibly non-deterministic.

23) A <derived table> or <lateral derived table> is *possibly non-deterministic* if the simply contained <query expression> is possibly non-deterministic.

24) A <table primary> is *possibly non-deterministic* if the simply contained <table name>, <query name>, <derived table>, <lateral derived table>, or <joined table> is possibly non-deterministic.

- 25) 14 A <table reference> is *possibly non-deterministic* if the simply contained <table primary> or <joined table> is possibly non-deterministic or if <sample clause> is specified.

Access Rules

- 1) If a <table primary> *TP* simply contains a <table or query name> that simply contains a <table name> *TN*, then:
 - a) Let *T* be the table identified by *TN*.
 - b) Case:
 - i) If *TN* is contained in a <search condition> immediately contained in an <assertion definition> or a <check constraint definition>, then the applicable privileges for the <authorization identifier> that owns the containing schema shall include REFERENCES on at least one column of *T*.
 - ii) Otherwise:
 - 1) Case:
 - A) If *TP* is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include SELECT on at least one column of *T*.
 - B) Otherwise, the current privileges shall include SELECT on at least one column of *T*.
 - 2) If *TP* simply contains <only spec> and *TN* identifies a typed table, then
 Case:
 - A) If *TP* is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges for the <authorization identifier> that owns the containing schema shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *T*.
 - B) Otherwise, the current privileges shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *T*.

General Rules

- 1) If a <table primary> *TP* simply contains a <table or query name> *TOQN*, then
 Case:
 - a) If *TOQN* simply contains a <query name> *QN*, then the result of *TP* is the table specified by *QN*.
 - b) If *TOQN* simply contains a <transition table name> *TTN*, then the result of *TP* is the table specified by *TTN*.

NOTE 153 — The table identified by a <transition table name> is a transition table as defined by the General Rules of Subclause 15.7, “Effect of deleting rows from base tables”, Subclause 15.10, “Effect of inserting tables into base tables”, or Subclause 15.13, “Effect of replacing rows in base tables”, as appropriate.
 - c) 09 Otherwise, let *T* be the table specified by the <table name> simply contained in *TP*.

Case:

- i) If ONLY is specified, then the result of *TP* is a table that consists of every row in *T*, except those rows that have a subrow in a proper subtable of *T*.
 - ii) Otherwise, the result of *TP* is a table that consists of every row of *T*.
- 2) 14 If a <derived table> or <lateral derived table> *LDT* simply containing <query expression> *QE* is specified, then the result of *LDT* is the result of *QE*.
- 3) If *TP* immediately contains a <joined table>, then the result of *TP* is the result of that <joined table>.
- 4) Let *TP* be the <table primary> immediately contained in a <table factor> *TF*. Let *RT* be the result of *TP*.

Case:

a) If <sample clause> is specified, then:

- i) Let *N* be the number of rows in *RT* and let *S* be the value of <sample percentage>.
- ii) If *S* is the null value or if $S < 0$ (zero) or if $S > 100$, then an exception condition is raised: *data exception — invalid sample size*.
- iii) If <repeatable clause> is specified, then let *RPT* be the value of <repeat argument>. If *RPT* is the null value, then an exception condition is raised: *data exception — invalid repeat argument in a sample clause*.
- iv) Case:
 - 1) If <sample method> specifies BERNOULLI, then the result of *TF* is a table containing approximately $(N*S/100)$ rows of *RT*. The probability of a row of *RT* being included in result of *TF* is $S/100$. Further, whether a given row of *RT* is included in result of *TF* is independent of whether other rows of *RT* are included in result of *TF*.
 - 2) Otherwise, result of *TF* is a table containing approximately $(N*S/100)$ rows of *RT*. The probability of a row of *RT* being included in result of *TF* is $S/100$.
- v) If *TF* contains outer references, then a table with identical rows is generated every time *TF* is evaluated with a given set of values for outer references.

NOTE 154 — “outer reference” is defined in Subclause 6.7, “<column reference>”.

b) Otherwise, result of *TF* is *RT*.

- 5) The result of a <table reference> *TR* is the result of immediately contained <table factor> or <joined table>.
- 6) Let *RV* be the range variable that is exposed by a <table factor> *TF*. The table associated with *RV* is the result of *TF*.

NOTE 155 — “table associated with *RV*” is defined in Subclause 4.14.6, “Operations involving tables”.

Conformance Rules

- 1) Without Feature S091, “Basic array support”, or Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <collection derived table>.
- 2) Without Feature T491, “LATERAL derived table”, conforming SQL language shall not contain a <lateral derived table>.

- 3) Without Feature T121, “WITH (excluding RECURSIVE) in query expression”, conforming SQL language shall not contain a <query name>.
- 4) Without Feature S111, “ONLY in query expressions”, conforming SQL language shall not contain a <table reference> that contains an <only spec>.
- 5) Without Feature F591, “Derived tables”, conforming SQL language shall not contain a <derived table>.
- 6) Without Feature T326, “Table functions”, conforming SQL language shall not contain a <table function derived table>.
- 7) Without Feature T613, “Sampling”, conforming SQL language shall not contain a <sample clause>.
- 8) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain a <transition table name>.
- 9) Without Feature S301, “Enhanced UNNEST”, in conforming SQL language, a <collection derived table> shall not simply contain more than one <collection value expression>.

7.7 <joined table>

Function

Specify a table derived from a Cartesian product, inner join, or outer join.

Format

```

<joined table> ::=
    <cross join>
  | <qualified join>
  | <natural join>

<cross join> ::=
    <table reference> CROSS JOIN <table factor>

<qualified join> ::=
    { <table reference> | <partitioned join table> }
    [ <join type> ] JOIN
    { <table reference> | <partitioned join table> }
    <join specification>

<partitioned join table> ::=
    <table factor> PARTITION BY
    <partitioned join column reference list>

<partitioned join column reference list> ::=
    <left paren> <partitioned join column reference>
    [ { <comma> <partitioned join column reference> }... ]
    <right paren>

<partitioned join column reference> ::=
    <column reference>

<natural join> ::=
    { <table reference> | <partitioned join table> }
    NATURAL [ <join type> ] JOIN
    { <table factor> | <partitioned join table> }

<join specification> ::=
    <join condition>
  | <named columns join>

<join condition> ::=
    ON <search condition>

<named columns join> ::=
    USING <left paren> <join column list> <right paren>

<join type> ::=
    INNER
  | <outer join type> [ OUTER ]

<outer join type> ::=
    LEFT
  | RIGHT

```

7.7 <joined table>

| FULL

```
<join column list> ::=
  <column name list>
```

Syntax Rules

- 1) Let *TRA* be the <table reference> or <table factor> that is the first operand of the <joined table>, and let *TRB* be the <table reference> or <table factor> that is the second operand of the <joined table>. Let *RTA* and *RTB* be the row types of *TRA* and *TRB*, respectively. Let *RVA* and *RVB* be the range variables of *TRA* and *TRB*, respectively. Let *CP* be:

```
SELECT *
FROM TRA, TRB
```

- 2) If *TRB* contains a <lateral derived table> containing an outer reference that references *TRA*, then <join type> shall not contain RIGHT or FULL.
- 3) If a <qualified join> or <natural join> is specified and a <join type> is not specified, then INNER is implicit.
- 4) If a <qualified join> containing a <join condition> is specified and a <value expression> directly contained in the <search condition> is a <set function specification>, then the <joined table> shall be contained in a <having clause> or <select list>, the <set function specification> shall contain an aggregated argument *AA* that contains an outer reference, and every column reference contained in *AA* shall be an outer reference.

NOTE 156 — “outer reference” is defined in Subclause 6.7, “<column reference>”.

- 5) The <search condition> shall not contain a <window function> without an intervening <query expression>.
- 6) If neither NATURAL is specified nor a <join specification> immediately containing a <named columns join> is specified, then the descriptors of the columns of the result of the <joined table> are the same as the descriptors of the columns of *CP*, with the possible exception of the nullability characteristics of the columns.
- 7) If a <partitioned join table> *PJT* is specified, then:
 - a) The <qualified join> or <natural join> shall specify an <outer join type>.
 - b) Each <partitioned join column reference> shall uniquely reference a column of the table referenced by the <table primary> simply contained in *PJT*. Such a column is called a *join partitioning column*.
 - c) If the first operand of the <qualified join> or <natural join> is a <partitioned join table>, then the <outer join type> shall be RIGHT or FULL.
 - d) If the second operand of the <qualified join> or <natural join> is a <partitioned join table>, then the <outer join type> shall be LEFT or FULL, and *TRB* shall not contain a <lateral derived table> containing an outer reference that references *TRA*.
- 8) If NATURAL is specified or if a <join specification> immediately containing a <named columns join> is specified, then:
 - a) If NATURAL is specified, then let *common column name* be a <field name> that is equivalent to the <field name> of exactly one field of *RTA* and the <field name> of exactly one field of *RTB*. *RTA* shall not have any duplicate common column names and *RTB* shall not have any duplicate common column names. Let *corresponding join columns* refer to all fields of *RTA* and *RTB* that have common column names, if any.

- b) If a <named columns join> is specified, then every <column name> in the <join column list> shall be equivalent to the <field name> of exactly one field of *RTA* and the <field name> of exactly one field of *RTB*. Let *common column name* be the name of such a column. Let *corresponding join columns* refer to the columns identified in the <join column list>.
- c) A corresponding join column shall not be a join partitioning column.
- d) Let *CA* and *CB* be a pair of corresponding join columns of *RTA* and *RTB*, respectively. *CA* and *CB* shall be comparable. *CA* and *CB* are operands of an equality operation, and the Syntax Rules of Subclause 9.9, “Equality operations”, apply.
- e) If there is at least one corresponding join column, then let *SLCC* be a <select list> of <derived column>s of the form

COALESCE (*RVA.C*, *RVB.C*) AS *C*

for every column *C* that is a corresponding join column, taken in order of their ordinal positions in *RTA*.

- f) If *RTA* contains at least one field that is not a corresponding join column, then let *SLTA* be a <select list> of <derived column>s of the form

RVA.C

for every field *C* of *RTA* that is not a corresponding join column, taken in order of their ordinal positions in *RTA*.

- g) If *RTB* contains at least one field that is not a corresponding join column, then let *SLTB* be a <select list> of <derived column>s of the form

RVB.C

for every field *C* of *RTB* that is not a corresponding join column, taken in order of their ordinal positions in *RTB*.

- h) Let the <select list> *SL* be defined as

Case:

- i) If all of the fields of *RTA* and *RTB* are corresponding join columns, then let *SL* be “*SLCC*”.
- ii) If *RTA* contains no corresponding join columns and *RTB* contains no corresponding join columns, then let *SL* be “*SLTA*, *SLTB*”.
- iii) If *RTA* contains no fields other than corresponding join columns, then let *SL* be “*SLCC*, *SLTB*”.
- iv) If *RTB* contains no fields other than corresponding join columns, then let *SL* be “*SLCC*, *SLTA*”.
- v) Otherwise, let *SL* be “*SLCC*, *SLTA*, *SLTB*”.

The descriptors of the columns of the result of the <joined table>, with the possible exception of the nullability characteristics of the columns, are the same as the descriptors of the columns of the result of

SELECT *SL* FROM *TRA*, *TRB*

- 9) A <joined table> is *possibly non-deterministic* if at least one of the following conditions is true:

7.7 <joined table>

- a) Either *TRA* or *TRB* is possibly non-deterministic.
 - b) A <join condition> that generally contains a possibly non-deterministic <value expression>, possibly non-deterministic <query specification>, or possibly non-deterministic <query expression> is specified.
 - c) NATURAL is specified, or a <join specification> immediately containing a <named columns join> is specified, and there is a common column name *CCN* such that the declared types of the two corresponding join columns identified by *CCN* have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.
 - d) The declared type of a join partitioning column is character string, user-defined type, TIME WITH TIME ZONE, or TIMESTAMP WITH TIME ZONE.
- 10) The declared type of the rows of the <joined table> is the row type *RT* defined by the sequence of (<field name>, <data type>) pairs indicated by the sequence of column descriptors of the <joined table> taken in order.
- 11) A column *CR* of the result of the <joined table> is *readily known not null* if all of the following conditions are true:
- a) The <joined table> does not simply contain an <outer join type>.
 - b) The <joined table> does not simply contain NATURAL.
 - c) The <joined table> does not simply contain a <partitioned join table>.
 - d) *CR* is a column reference that references a column that is readily known not null.
- 12) A column *CR* of the result of the <joined table> is *known not null* if any of the following conditions is true:
- a) *CR* is readily known not null.
 - b) If the SQL-implementation supports Feature T101, “Enhanced nullability determination”, then *CR* is known not null if *CR* is not possibly nullable according to the following conditions:
 - i) For every column *CR* of the result of the <joined table> that corresponds to a field *CA* of *RTA* that is not a corresponding join column or a join partitioning column, *CR* is *possibly nullable* if any of the following conditions are true:
 - 1) RIGHT or FULL is specified.
 - 2) INNER, LEFT, or CROSS JOIN is specified or implicit and *CA* is possibly nullable.
 - ii) For every column *CR* of the result of the <joined table> that corresponds to a field *CB* of *RTB* that is not a corresponding join column or a join partitioning column, *CR* is *possibly nullable* if any of the following conditions are true:
 - 1) LEFT or FULL is specified.
 - 2) INNER, RIGHT, or CROSS JOIN is specified or implicit and *CB* is possibly nullable.
 - iii) For every column *CR* of the result of the <joined table> that corresponds to a corresponding join column *CA* of *RTA* and a corresponding join column *CB* of *RTB*, *CR* is *possibly nullable* if any of the following conditions are true:
 - 1) LEFT or FULL is specified and *CA* is possibly nullable, or
 - 2) RIGHT or FULL is specified and *CB* is possibly nullable.

- iv) A column *CR* of the result of the <joined table> that corresponds to a join partitioning column *JPC* is *possibly nullable* if *JPC* is possibly nullable.
- c) If *CR* conforms to an implementation-defined rule that correctly infers that the value of *CR* cannot be null.

Access Rules

None.

General Rules

1) Case:

- a) If a <cross join> is specified, then let *T* be *CP*.
- b) If a <join condition> is specified, then let *SC* be the <search condition> and let *T* be

CP
WHERE *SC*

- c) If NATURAL is specified or <named columns join> is specified, then

Case:

- i) If there are corresponding join columns, then let *N* be the number of corresponding join columns and let *CJCN_i*, 1 (one) ≤ *i* ≤ *N*, be the field name of the *i*-th corresponding join column, and let *T* be

CP
WHERE *RVA.CJCN₁* = *RVB.CJCN₁*
AND ...
AND *RVA.CJCN_N* = *RVB.CJCN_N*

- ii) Otherwise, let *T* be *CP*.

- 2) Let *TR* be the result of evaluating *T*, let *DA* and *DB* be the degrees of *TRA* and *TRB*, respectively, and let *TVA* be the result of evaluating *TRA*.
- 3) If LEFT is specified and the second operand does not specify <partitioned join table>, then:
 - a) Let *PA* be the collection of rows of *TVA* for which there exists some row *R* in *TR* and some row *RA* in *TVA* such that the values of the first *DA* fields of *R* are identical to the values of the corresponding fields of *RA*.
 - b) Let *UA* be those rows of *TVA* that are not in *PA*.
 - c) Let *XA* be *UA* extended on the right with *DB* columns containing the null value.
 - d) Let *XNA* be an effective distinct name for *XA*. Let *TN* be an effective name for *T*.
- 4) If RIGHT or FULL is specified or if LEFT is specified and the second operand specifies <partitioned join table>, then:

- a) Let *TVB* be the result of evaluating *TRB*.

NOTE 157 — It follows from the Syntax Rules that *TRB* does not contain a <lateral derived table> containing an outer reference that references *TRA*. This ensures that it is possible to evaluate *TRB* in isolation.

- b) Case:

- i) If the first operand specifies <partitioned join table>, then

Case:

- 1) If *TVA* is empty, then let *NA* be 0 (zero).
- 2) Otherwise, *TVA* is partitioned into the minimum numbers of partitions such that for each join partitioning column *JPC* of each partition, no two values of *JPC* are distinct. If the declared type of a join partitioning column is a user-defined type and the comparison of that column results in *Unknown* for two rows of *TVA*, then the assignment of those rows to partitions is implementation-dependent. Let *NA* be the number of partitions. Let GA_1, \dots, GA_{NA} be an enumeration of the partitions.

- ii) Otherwise, let *NA* be 1 (one), and let GA_1 be *TVA*.

- c) Case:

- i) If the second operand specifies <partitioned join table>, then

Case:

- 1) If *TVB* is empty, then let *NB* be 0 (zero).
- 2) Otherwise, *TVB* is partitioned into the minimum numbers of partitions such that for each join partitioning column *JPC* of each partition, no two values of the join partitioning column are distinct. If the declared type of a join partitioning column is a user-defined type and the comparison of that column results in *Unknown* for two rows of *TVB*, then the assignment of those rows to partitions is implementation-dependent. Let *NB* be the number of partitions. Let GB_1, \dots, GB_{NB} be an enumeration of the partitions.

- ii) Otherwise, let *NB* be 1 (one), and let GB_1 be *TVB*.

- d) For each *i* between 1 (one) and *NA*, and for each *j* between 1 (one) and *NB*, let $PA_{i,j}$ be the collection of rows *RA* of GA_i for which there exists a row *RB* in GB_j such that the concatenation of *RA* and *RB* is in *TR*.
- e) For each *i* between 1 (one) and *NA*, and for each *j* between 1 (one) and *NB*, let $PB_{i,j}$ be the collection of rows *RB* of GB_j for which there exists a row *RA* in GA_i such that the concatenation of *RA* and *RB* is in *TR*.
- f) For each *i* between 1 (one) and *NA*, and for each *j* between 1 (one) and *NB*, let $UA_{i,j}$ be the collection of rows of GA_i that are not in $PA_{i,j}$.
- g) For each *i* between 1 (one) and *NA*, and for each *j* between 1 (one) and *NB*, let $UB_{i,j}$ be the collection of rows of GB_j that are not in $PB_{i,j}$.
- h) For each *i* between 1 (one) and *NA*, and for each *j* between 1 (one) and *NB*, let $XA_{i,j}$ be $UA_{i,j}$ extended on the right with *DB* columns, with declared types and values determined as follows. For each *k* between

1 (one) and DB , the declared type of the $(DA + k)$ -th column is the declared type of the k -th column of TVB , and the value is

Case:

- i) If the k -th column of TVB is a join partitioning column, then the common value of the k -th column of GB_j .
 - ii) Otherwise, the null value.
- i) For each i between 1 (one) and NA , and for each j between 1 (one) and NB , let $XB_{i,j}$ be $UB_{i,j}$ extended on the left with DA columns, with declared types and values determined as follows. For each k between 1 (one) and DA , the declared type of the k -th column is the declared type of the k -th column of TVA , and the value is

Case:

- i) If the k -th column of TVA is a join partitioning column, then the common value of the k -th column of GA_i .
 - ii) Otherwise, the null value.
- j) Let XA be the collection of all rows in $XA_{i,j}$ for all i between 1 (one) and NA and all j between 1 (one) and NB .

NOTE 158 — If NA is 0 (zero), then XA is empty.

- k) Let XB be the collection of all rows in $XB_{i,j}$ for all i between 1 (one) and NA and all j between 1 (one) and NB .

NOTE 159 — If NB is 0 (zero), then XB is empty.

- l) Let XNB be an effective distinct name for XB .

5) Case:

- a) If INNER or <cross join> is specified, then let S be TR .
- b) If LEFT is specified, then let S be the result of:

```
SELECT *
FROM TN
UNION ALL
SELECT *
FROM XNA
```

- c) If RIGHT is specified, then let S be the result of:

```
SELECT *
FROM TN
UNION ALL
SELECT *
FROM XNB
```

- d) If FULL is specified, then let S be the result of:

```
SELECT *
FROM TN
```

7.7 <joined table>

```
UNION ALL
SELECT *
FROM XNA
UNION ALL
SELECT *
FROM XNB
```

6) Let *SN* be an effective name of *S*.

Case:

a) If NATURAL is specified or a <named columns join> is specified, then:

i) Let CS_i be a name for the *i*-th column of *S*. Column CS_i of *S* corresponds to the *i*-th field of *RTA* if *i* is less than or equal to *DA*. Column CS_j of *S* corresponds to the (*j*−*DA*)-th field of *RTB* for *j* greater than *DA*.

ii) If there is at least one corresponding join column, then let *SLCC* be a <select list> of derived columns of the form

$COALESCE (CS_i, CS_j)$

for every pair of columns CS_i and CS_j , where CS_i and CS_j correspond to fields of *RTA* and *RTB* that are a pair of corresponding join columns.

iii) If *RTA* contains one or more fields that are not corresponding join columns, then let *SLTA* be a <select list> of the form:

CS_i

for every column CS_i of *S* that corresponds to a field of *RTA* that is not a corresponding join column, taken in order of their ordinal position in *S*.

iv) If *RTB* contains one or more fields that are not corresponding join columns, then let *SLTB* be a <select list> of the form:

CS_j

for every column CS_j of *S* that corresponds to a field of *RTB* that is not a corresponding join column, taken in order of their ordinal position in *S*.

v) Let the <select list> *SL* be defined as

Case:

1) If all the fields of *RTA* and *RTB* are corresponding join columns, then let *SL* be

SLCC

2) If *RTA* contains no corresponding join columns and *RTB* contains no corresponding join columns, then let *SL* be

SLTA, SLTB

3) If *RTA* contains no fields other than corresponding join columns, then let *SL* be

SLCC , *SLTB*

- 4) If *RTB* contains no fields other than corresponding join columns, then let *SL* be

SLCC , *SLTA*

- 5) Otherwise, let *SL* be

SLCC , *SLTA* , *SLTB*

- vi) The result of the <joined table> is the result of:

```
SELECT SL
FROM SN
```

- b) Otherwise, the result of the <joined table> is *S*.

Conformance Rules

- 1) Without Feature F401, “Extended joined table”, conforming SQL language shall not contain a <cross join>.
- 2) Without Feature F401, “Extended joined table”, conforming SQL language shall not contain a <natural join>.
- 3) Without Feature F401, “Extended joined table”, conforming SQL language shall not contain FULL.
- 4) Without Feature F402, “Named column joins for LOBs, arrays, and multisets”, conforming SQL language shall not contain a <joined table> that simply contains either <natural join> or <named columns join> in which, if *C* is a corresponding join column, the declared type of *C* is LOB-ordered, array-ordered, or multiset-ordered.

NOTE 160 — If *C* is a corresponding join column, then the Conformance Rules of [Subclause 9.9](#), “Equality operations”, also apply.

- 5) Without Feature F403, “Partitioned join tables”, conforming SQL language shall not contain <partitioned join table>.

7.8 <where clause>

Function

Specify a table derived by the application of a <search condition> to the result of the preceding <from clause>.

Format

```
<where clause> ::=
  WHERE <search condition>
```

Syntax Rules

- 1) If a <value expression> directly contained in the <search condition> is a <set function specification>, then the <where clause> shall be contained in a <having clause> or <select list>, the <set function specification> shall contain a column reference, and every column reference contained in an aggregated argument of the <set function specification> shall be an outer reference.

NOTE 161 — *outer reference* is defined in Subclause 6.7, “<column reference>”.

- 2) The <search condition> shall not contain a <window function> without an intervening <query expression>.

Access Rules

None.

General Rules

- 1) Let *T* be the result of the preceding <from clause>.
- 2) The <search condition> is effectively evaluated for each row of *T*. The result of the <where clause> is a table of those rows of *T* for which the result of the <search condition> is True.

Conformance Rules

- 1) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <value expression> directly contained in a <where clause> that contains a <column reference> that references a <derived column> that generally contains a <set function specification> without an intervening <routine invocation>.

7.9 <group by clause>

Function

Specify a grouped table derived by the application of the <group by clause> to the result of the previously specified clause.

Format

```

<group by clause> ::=
    GROUP BY [ <set quantifier> ] <grouping element list>

<grouping element list> ::=
    <grouping element> [ { <comma> <grouping element> }... ]

<grouping element> ::=
    <ordinary grouping set>
    | <rollup list>
    | <cube list>
    | <grouping sets specification>
    | <empty grouping set>

<ordinary grouping set> ::=
    <grouping column reference>
    | <left paren> <grouping column reference list> <right paren>

<grouping column reference> ::=
    <column reference> [ <collate clause> ]

<grouping column reference list> ::=
    <grouping column reference> [ { <comma> <grouping column reference> }... ]

<rollup list> ::=
    ROLLUP <left paren> <ordinary grouping set list> <right paren>

<ordinary grouping set list> ::=
    <ordinary grouping set> [ { <comma> <ordinary grouping set> }... ]

<cube list> ::=
    CUBE <left paren> <ordinary grouping set list> <right paren>

<grouping sets specification> ::=
    GROUPING SETS <left paren> <grouping set list> <right paren>

<grouping set list> ::=
    <grouping set> [ { <comma> <grouping set> }... ]

<grouping set> ::=
    <ordinary grouping set>
    | <rollup list>
    | <cube list>
    | <grouping sets specification>
    | <empty grouping set>

<empty grouping set> ::=
    <left paren> <right paren>

```

Syntax Rules

- 1) Each <grouping column reference> shall unambiguously reference a column of the table resulting from the <from clause>. A column referenced in a <group by clause> is a *grouping column*.

NOTE 162 — “Column reference” is defined in Subclause 6.7, “<column reference>”.

- 2) Each <grouping column reference> is an operand of a grouping operation. The Syntax Rules of Subclause 9.10, “Grouping operations”, apply.
- 3) For every <grouping column reference> *GC*,
Case:
 - a) If <collate clause> is specified, then let *CS* be the collation identified by <collation name>. The declared type of the column reference shall be character string. The declared type of *GC* is that of its column reference, except that *CS* is the declared type collation and the collation derivation is *explicit*.
 - b) Otherwise, the declared type of *GC* is the declared type of its column reference.
- 4) Let *QS* be the <query specification> that simply contains the <group by clause>, and let *SL*, *FC*, *WC*, *GBC*, and *HC* be the <select list>, the <from clause>, the <where clause> if any, the <group by clause>, and the <having clause> if any, respectively, that are simply contained in *QS*.
- 5) Let *QSSQ* be the explicit or implicit <set quantifier> immediately contained in *QS*.
- 6) Let *GBSQ* be the <set quantifier> immediately contained in <group by clause>, if any; otherwise, let *GBSQ* be ALL.
- 7) Let *SLI* be obtained from *SL* by replacing every <asterisk> and <asterisked identifier chain> using the syntactic transformations in the Syntax Rules of Subclause 7.12, “<query specification>”.
- 8) A <group by clause> is *primitive* if it does not contain a <rollup list>, <cube list>, <grouping sets specification>, or <grouping column reference list>, and does not contain both a <grouping column reference> and an <empty grouping set>.
- 9) A <group by clause> is *simple* if it does not contain a <rollup list>, <cube list> or <grouping sets specification>.
- 10) If *GBC* is a simple <group by clause> that is not primitive, then *GBC* is transformed into a primitive <group by clause> as follows:
 - a) Let *NSGB* be the number of <grouping column reference>s contained in *GBC*.
 - b) Case:
 - i) If *NSGB* is 0 (zero), then *GBC* is replaced by

GROUP BY ()
 - ii) Otherwise:
 - 1) Let *SGCR*₁, ... *SGCR*_{*NSGB*} be an enumeration of the <grouping column reference>s contained in *GBC*.
 - 2) *GBC* is replaced by

```
GROUP BY  $SGCR_1$ , ...
 $SGCR_{NSGB}$ 
```

NOTE 163 — That is, a simple <group by clause> that is not primitive may be transformed into a primitive <group by clause> by deleting all parentheses, and deleting extra <comma>s as necessary for correct syntax. If there are no grouping columns at all (for example, GROUP BY (), ()), this is transformed to the canonical form GROUP BY ().

- 11) If GBC is a primitive <group by clause>, then let $SLNEW$ and $HCNEW$ be obtained from SLI and HC , respectively, by replacing every <grouping operation> by the exact numeric literal 0 (zero). QS is equivalent to:

```
SELECT  $QSSQ$   $SLNEW$   $FC$ 
WC  $GBC$   $HCNEW$ 
```

- 12) Let $OGSL$ be an <ordinary grouping set list>. The *concatenation* of $OGSL$ is defined as follows:

- Let $NGCR$ be the number of <grouping column reference>s simply contained in $OGSL$ and let GCR_j , $1 \text{ (one)} \leq j \leq NGCR$, be an enumeration of those <grouping column reference>s, in order from left to right.
- The concatenation of $OGSL$ is the <ordinary grouping set list>

```
 $GCR_1$ , ...,  $GCR_{NGCR}$ 
```

NOTE 164 — Thus, the concatenation of $OGSL$ may be formed by erasing all parentheses. For example, the concatenation of “(A, B), (C, D)” is “A, B, C, D”.

- 13) Let RL be a <rollup list>. Let OGS_i range over the n <ordinary grouping set>s contained in RL .

- For each i between 1 (one) and n , let $COGS_i$ be the concatenation of the <ordinary grouping set list>

```
 $OGS_1$ ,  $OGS_2$ , ...,  $OGS_i$ 
```

- RL is equivalent to:

```
GROUPING SETS (
  (  $COGS_n$  ),
  (  $COGS_{n-1}$  ),
  (  $COGS_{n-2}$  ),
  ...
  (  $COGS_1$  ),
  ( ) )
```

NOTE 165 — The result of the transform is to replace RL with a <grouping sets specification> that contains a <grouping set> for every initial sublist of the <ordinary grouping set list> of the <rollup list>, obtained by dropping <ordinary grouping set>s from the right, one by one, and concatenating each <ordinary grouping set list> so obtained. The <empty grouping set> is regarded as the shortest such initial sublist. For example, “ROLLUP ((A, B), (C, D))” is equivalent to “GROUPING SETS ((A, B, C, D), (A, B), ())”.

- 14) Let CL be a <cube list>. Let OGS_i range over the n <ordinary grouping set>s contained in CL . CL is transformed as follows:

- Let $M = 2^n - 1$ (one).
- For each i between 1 (one) and M :

7.9 <group by clause>

- i) Let BSL_i be the binary number consisting of n bits (binary digits) whose value is i .
- ii) For each j between 1 (one) and n , let $B_{i,j}$ be the j -th bit, counting from left to right, in BSL_i .
- iii) For each j between 1 (one) and n , let $GSLCR_{i,j}$ be

Case:

- 1) If $B_{i,j}$ is 0 (zero), then the zero-length string.
- 2) If $B_{i,j}$ is 1 (one) and $B_{i,k}$ is 0 (zero) for all $k < j$, then OGS_j .
- 3) Otherwise, <comma> followed by OGS_j .

- iv) Let GSL_i be the concatenation of the <ordinary grouping set list>

$GSLCR_{i,1} \text{ } GSLCR_{i,2} \text{ } \dots \text{ } GSLCR_{i,n}$

- c) CL is equivalent to

GROUPING SETS ((GSL_M), (GSL_{M-1}), . . . , (GSL_1), ())

NOTE 166 — The result of the transform is to replace CL with a <grouping sets specification> that contains a <grouping set> for all possible subsets of the set of <ordinary grouping set>s in the <ordinary grouping set list> of the <cube list>, including <empty grouping set> as the empty subset with no <ordinary grouping set>s.

For example, CUBE (A , B , C) is equivalent to:

```
GROUPING SETS ( /*  $BSL_i$  */
  ( A , B , C ) , /* 111 */
  ( A , B ) , /* 110 */
  ( A , C ) , /* 101 */
  ( A ) , /* 100 */
  ( B , C ) , /* 011 */
  ( B ) , /* 010 */
  ( C ) , /* 001 */
  ( )
)
```

As another example, CUBE ((A , B) , (C , D)) is equivalent to:

```
GROUPING SETS ( /*  $BSL_i$  */
  ( A , B , C , D ) , /* 11 */
  ( A , B ) , /* 10 */
  ( C , D ) , /* 01 */
  ( )
)
```

- 15) Let $GSSA$ be a <grouping sets specification>. If $GSSA$ simply contains another <grouping sets specification> $GSSB$, then $GSSA$ is transformed as follows:

- a) Let NA be the number of <grouping set>s simply contained in $GSSA$, and let NB be the number of <grouping set>s simply contained in $GSSB$.
- b) Let GSA_i be an enumeration of the <grouping set>s simply contained in $GSSA$, for $1 \text{ (one)} \leq i \leq NA$.

- c) Let GSB_i be an enumeration of the <grouping set>s simply contained in $GSSB$, $1 \text{ (one)} \leq i \leq NB$.
- d) Let k be the value such that $GSSB = GSA_k$.
- e) $GSSA$ is equivalent to

GROUPING SETS (
 $GSA_1, GSA_2, \dots, GSA_{k-1},$
 $GSB_1, \dots, GSB_{NB},$
 $GSA_{k+1}, \dots, GSA_{NA})$

NOTE 167 — Thus, the nested <grouping sets specification> is removed by simply “promoting” each of its <grouping set>s to be a <grouping set> of the encompassing <grouping sets specification>.

16) Let CGB be a <group by clause> that is not simple. CGB is transformed as follows:

- a) The preceding Syntax Rules are applied repeatedly to eliminate any <grouping sets specification> that is nested in another <grouping sets specification>, as well as any <rollup list> and any <cube list>.

NOTE 168 — As a result, CGB is either a single <grouping sets specification> or a list of two or more <grouping set>s, each of which is an <ordinary grouping set>, an <empty grouping set>, or a <grouping sets specification> that contains only <ordinary grouping set>s and <empty grouping set>s. There are no remaining <rollup list>s, <cube list>s, or nested <grouping sets specification>s.

- b) Any <grouping element> GS that is an <ordinary grouping set> or an <empty grouping set> is replaced by the <grouping sets specification>

GROUPING SETS (GS)

NOTE 169 — As a result, CGB is a list of one or more <grouping sets specification>s.

- c) If CGB contains two or more <grouping sets specification>s, then let $GSSX$ and $GSSY$ be the first two <grouping sets specification>s in CGB . CGB is transformed by replacing “ $GSSX$ <comma> $GSSY$ ” as follows:

- i) Let NX be the number of <grouping set>s in $GSSX$ and let NY be the number of <grouping set>s in $GSSY$.
- ii) Let GSX_i , $1 \text{ (one)} \leq i \leq NX$, be the <grouping set>s contained in $GSSX$, and let GSY_i , $1 \text{ (one)} \leq i \leq NY$, be the <grouping set>s contained in $GSSY$.
- iii) Let $MX(i)$ be the number of <grouping column reference>s in GSX_i , and let $MY(i)$ be the number of <grouping column reference>s in GSY_i .

NOTE 170 — If GSX_i is <empty grouping set>, then $MX(i)$ is 0 (zero); and similarly for GSY_i .

- iv) Let $GCRX_{i,j}$, $1 \text{ (one)} \leq j \leq MX(i)$ be the <grouping column reference>s contained in GSX_i , and let $GCRY_{i,j}$, $1 \text{ (one)} \leq j \leq MY(i)$ be the <grouping column reference>s contained in GSY_i .

NOTE 171 — If GSX_i is <empty grouping set>, then there are no $GCRX_{i,j}$; and similarly for GSY_i .

- v) For each a between 1 (one) and NX and each b between 1 (one) and NY , let $GST_{a,b}$ be

($GCRX_{a,1}, \dots, GCRX_{a,MX(a)}, GCRY_{b,1}, \dots, GCRY_{b,MY(b)})$

that is, an <ordinary grouping set> consisting of $GCRX_{a,j}$ for all j between 1 (one) and $MX(a)$, followed by $GCRY_{b,j}$ for all j between 1 (one) and $MY(b)$.

- vi) *CGB* is transformed by replacing “*GSSX* <comma> *GSSY*” with

```
GROUPING SETS (
  GST1,1, . . . , GST1,NY,
  GST2,1, . . . , GST2,NY,
  . . .
  GSTNX,1 , . . . , GSTNX,NY
)
```

NOTE 172 — Thus each <ordinary grouping set> in *GSSA* is “concatenated” with each <ordinary grouping set> in *GSSB*. For example,

```
GROUP BY GROUPING SETS ((A, B), (C)),
      GROUPING SETS ((X, Y), ( ))
```

is transformed to

```
GROUP BY GROUPING SETS ((A, B, X, Y), (A, B),
      (C, X, Y), (C))
```

- d) The previous subrule of this Syntax Rule is applied repeatedly until *CGB* consists of a single <grouping sets specification>.
- 17) If <grouping element list> consists of a single <grouping sets specification> *GSS* that contains only <ordinary grouping set>s or <empty grouping set>s, then:
- Let *m* be the number of <grouping set>s contained in *GSS*.
 - Let *GS_i*, $1 \leq i \leq m$, range over the <grouping set>s contained in *GSS*.
 - Let *p* be the number of distinct <column reference>s that are contained in *GSS*.
 - Let *PC* be an ordered list of these <column reference>s ordered according to their left-to-right occurrence in the list.
 - Let *PC_k*, $1 \leq k \leq p$, be the *k*-th <column reference> in *PC*.
 - Let *DTPC_k* be the declared type of the column identified by *PC_k*.
 - Let *NDC* be the number of <derived column>s simply contained in *SLI*.
 - Let *DC_q*, $1 \leq q \leq NDC$, be an enumeration of the <derived column>s simply contained in *SLI*, in order from left to right.
 - Let *DCN_q* be the column name of *DC_q*, $1 \text{ (one)} \leq q \leq NDC$.
 - Let *VE_q*, $1 \text{ (one)} \leq q \leq NDC$, be the <value expression> simply contained in *DC_q*.
 - Let *XN_k*, $1 \text{ (one)} \leq k \leq p$, *YN_k*, $1 \text{ (one)} \leq k \leq p$, and *ZN_q*, $1 \text{ (one)} \leq q \leq NDC$, be implementation-dependent column names that are all distinct from one another.
 - Let *SL2* be the <select list>:

```
PC1 AS XN1, GROUPING (PC1) AS YN1,
. . . ,
```


PC_p AS XN_p , GROUPING (PC_p) AS YN_p ,
 VE_1 AS ZN_1 , ..., VE_{NDC} AS ZN_{NDC}

m) For each GS_i :

i) If GS_i is an <empty grouping set>, then let $n(i)$ be 0 (zero). If GS_i is a <grouping column reference>, then let $n(i)$ be 1 (one). Otherwise, let $n(i)$ be the number of <grouping column reference>s contained in the <grouping column reference list>.

ii) Let $GCR_{i,j}$, $1 \leq j \leq n(i)$, range over the <grouping column reference>s contained in GS_i .

iii) Case:

1) If GS_i is an <ordinary grouping set>, then

A) Transform $SL2$ to obtain $SL3$, and transform HC to obtain $HC3$, as follows:

For every PC_k , if there is no j such that $PC_k = GCR_{i,j}$, then make the following replacements in $SL2$ and HC :

I) Replace each <grouping operation> in $SL2$ and HC that contains a <column reference> that references PC_k by the <literal> 1 (one).

II) Replace each <column reference> in $SL2$ and HC that references PC_k by

CAST (NULL AS $DTPC_k$)

B) Transform $SL3$ to obtain $SLNEW$, and transform $HC3$ to obtain $HCNEW$ by replacing each <grouping operation> that remains in $SL3$ and $HC3$ by the <literal> 0 (zero).

NOTE 173 — Thus the value of a <grouping operation> is 0 (zero) if the grouping column referenced by the <grouping operation> is among the $GCR_{i,j}$ and 1(one) if it is not.

C) Let $GSSQL_i$ be:

```
SELECT QSSQ SLNEW
FC
WC
GROUP BY GCRi,1, ..., GCRi,n(i)
HCNEW
```

2) If GS_i is an <empty grouping set>, then

A) Transform $SL2$ to obtain $SLNEW$, and transform HC to obtain $HCNEW$, as follows:

For every k , $1 \leq k \leq p$:

I) Replace each <grouping operation> in $SL2$ and HC that contains a <column reference> that references PC_k by the <literal> 1 (one).

II) Replace each <column reference> in $SL2$ and HC that references PC_k by

CAST (NULL AS $DTPC_k$)

B) Let $GSSQL_i$ be

7.9 <group by clause>

```

SELECT QSSQ SLNEW
FC
WC
GROUP BY ( )
HCNEW

```

n) Let GU be:

```

GSSQL1
  UNION GBSQ
GSSQL2
  UNION GBSQ
...
  UNION GBSQ
GSSQLm

```

o) QS is equivalent to

```

SELECT QSSQ ZN1 AS DC1, ..., ZNNDC AS DCNDC
FROM ( GU )

```

Access Rules

None.

General Rules

NOTE 174 — As a result of the syntactic transformations specified in the Syntax Rules of this Subclause, only primitive <group by clause>s are left to consider.

- 1) If no <where clause> is specified, then let T be the result of the preceding <from clause>; otherwise, let T be the result of the preceding <where clause>.
- 2) Case:
 - a) If there are no grouping columns, then the result of the <group by clause> is the grouped table consisting of T as its only group.
 - b) Otherwise, the result of the <group by clause> is a partitioning of the rows of T into the minimum number of groups such that, for each grouping column of each group, no two values of that grouping column are distinct. If the declared type of a grouping column is a user-defined type and the comparison of that column results in *Unknown* for two rows of T , then the assignment of those rows to groups in the result of the <group by clause> is implementation-dependent.
- 3) When a <search condition> or <value expression> is applied to a group, a reference CR to a column that is functionally dependent on the grouping columns is understood as follows.

Case:

- a) If CR is a group-invariant column reference, then it is a reference to the common value in that column of the rows in that group. If the most specific type of the column is character, datetime with time zone,

or a user-defined type, then the value is an implementation-dependent value that is not distinct from the value of the column in each row of the group.

- b) Otherwise, *CR* is a within-group-varying column reference, and as such, it is a reference to the value of the column in each row of a given group determined by the grouping columns, to be used to construct the argument source of a <set function specification>.

Conformance Rules

- 1) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <rollup list>.
- 2) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <cube list>.
- 3) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <grouping sets specification>.
- 4) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain an <empty grouping set>.
- 5) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain an <ordinary grouping set> that contains a <grouping column reference list>.
- 6) Without Feature T432, “Nested and concatenated GROUPING SETS”, conforming SQL language shall not contain a <grouping set list> that contains a <grouping sets specification>.
- 7) Without Feature T432, “Nested and concatenated GROUPING SETS”, conforming SQL language shall not contain a <group by clause> that simply contains a <grouping sets specification> *GSS* where *GSS* is not the only <grouping element> simply contained in the <group by clause>.

NOTE 175 — The Conformance Rules of [Subclause 9.10, “Grouping operations”](#), also apply.

- 8) Without Feature T434, “GROUP BY DISTINCT”, conforming SQL language shall not contain a <group by clause> that simply contains a <set quantifier>.

7.10 <having clause>

Function

Specify a grouped table derived by the elimination of groups that do not satisfy a <search condition>.

Format

```
<having clause> ::=
    HAVING <search condition>
```

Syntax Rules

- 1) Let *HC* be the <having clause>. Let *TE* be the <table expression> that immediately contains *HC*. If *TE* does not immediately contain a <group by clause>, then “GROUP BY ()” is implicit. Let *T* be the descriptor of the table defined by the <group by clause> *GBC* immediately contained in *TE* and let *R* be the result of *GBC*.
 - 2) Let *G* be the set consisting of every column referenced by a <column reference> contained in *GBC*.
 - 3) Each column reference directly contained in the <search condition> shall be one of the following:
 - a) An unambiguous reference to a column that is functionally dependent on *G*.
 - b) An outer reference.
- NOTE 176 — See also the Syntax Rules of Subclause 6.7, “<column reference>”.
- 4) The <search condition> shall not contain a <window function> without an intervening <query expression>.
 - 5) The row type of the result of the <having clause> is the row type *RT* of *T*.

Access Rules

None.

General Rules

- 1) The <search condition> is evaluated for each group of *R*. The result of the <having clause> is a grouped table of those groups of *R* for which the result of the <search condition> is True.

Conformance Rules

- 1) Without Feature T301, “Functional dependencies”, in conforming SQL language, each column reference directly contained in the <search condition> shall be one of the following:
 - a) An unambiguous reference to a grouping column of *T*.
 - b) An outer reference.

- 2) Without Feature T301, “Functional dependencies”, in conforming SQL language, each column reference contained in a <query expression> in the <search condition> that references a column of *T* shall be one of the following:
 - a) An unambiguous reference to a grouping column of *T*.
 - b) Contained in an aggregated argument of a <set function specification>.

7.11 <window clause>

Function

Specify one or more window definitions.

Format

```

<window clause> ::=
    WINDOW <window definition list>

<window definition list> ::=
    <window definition> [ { <comma> <window definition> }... ]

<window definition> ::=
    <new window name> AS <window specification>

<new window name> ::=
    <window name>

<window specification> ::=
    <left paren> <window specification details> <right paren>

<window specification details> ::=
    [ <existing window name> ]
    [ <window partition clause> ]
    [ <window order clause> ]
    [ <window frame clause> ]

<existing window name> ::=
    <window name>

<window partition clause> ::=
    PARTITION BY <window partition column reference list>

<window partition column reference list> ::=
    <window partition column reference>
    [ { <comma> <window partition column reference> }... ]

<window partition column reference> ::=
    <column reference> [ <collate clause> ]

<window order clause> ::=
    ORDER BY <sort specification list>

<window frame clause> ::=
    <window frame units> <window frame extent>
    [ <window frame exclusion> ]

<window frame units> ::=
    ROWS
    | RANGE

<window frame extent> ::=
    <window frame start>
    | <window frame between>

```

```
<window frame start> ::=
    UNBOUNDED PRECEDING
  | <window frame preceding>
  | CURRENT ROW

<window frame preceding> ::=
    <unsigned value specification> PRECEDING

<window frame between> ::=
    BETWEEN <window frame bound 1> AND <window frame bound 2>

<window frame bound 1> ::=
    <window frame bound>

<window frame bound 2> ::=
    <window frame bound>

<window frame bound> ::=
    <window frame start>
  | UNBOUNDED FOLLOWING
  | <window frame following>

<window frame following> ::=
    <unsigned value specification> FOLLOWING

<window frame exclusion> ::=
    EXCLUDE CURRENT ROW
  | EXCLUDE GROUP
  | EXCLUDE TIES
  | EXCLUDE NO OTHERS
```

Syntax Rules

- 1) Let *TE* be the <table expression> that immediately contains the <window clause>.
- 2) <new window name> *NWN1* shall not be contained in the scope of another <new window name> *NWN2* such that *NWN1* and *NWN2* are equivalent.
- 3) Let *WDEF* be a <window definition>.
- 4) Each <column reference> contained in the <window partition clause> or <window order clause> of *WDEF* shall unambiguously reference a column of the derived table *T* that is the result of *TE*. A column referenced in a <window partition clause> is a *partitioning column*. Each partitioning column is an operand of a grouping operation, and the Syntax Rules of [Subclause 9.10, “Grouping operations”](#), apply.

NOTE 177 — If *T* is a grouped table, then the <column reference>s contained in <window partition clause> or <window order clause> shall reference columns of the grouped table obtained by performing the syntactic transformation in [Subclause 7.12, “<query specification>”](#).

- 5) For every <window partition column reference> *PC*,

Case:

- a) If <collate clause> is specified, then let *CS* be the collation identified by <collation name>. The declared type of the column reference shall be character string. The declared type of *PC* is that of its column reference, except that *CS* is the declared type collation and the collation derivation is *explicit*.
- b) Otherwise, the declared type of *PC* is the declared type of its column reference.

7.11 <window clause>

- 6) If *T* is a grouped table, then let *G* be the set of grouping columns of *T*. Each column reference contained in <window clause> that references a column of *T* shall reference a column that is functionally dependent on *G* or be contained in an aggregated argument of a <set function specification>.
- 7) A <window clause> shall not contain a <window function> without an intervening <query expression>.
- 8) If *WDEF* specifies <window frame between>, then:
 - a) <window frame bound 1> shall not specify UNBOUNDED FOLLOWING.
 - b) <window frame bound 2> shall not specify UNBOUNDED PRECEDING.
 - c) If <window frame bound 1> specifies CURRENT ROW, then <window frame bound 2> shall not specify <window frame preceding>.
 - d) If <window frame bound 1> specifies <window frame following>, then <window frame bound 2> shall not specify <window frame preceding> or CURRENT ROW.
- 9) If *WDEF* specifies <window frame extent>, and does not specify <window frame between>, then let *WAGS* be the <window frame start>. The <window frame extent> is equivalent to
 BETWEEN *WAGS* AND CURRENT ROW
- 10) If *WDEF* specifies an <existing window name> *EWN*, then:
 - a) *WDEF* shall be within the scope of a <window name> that is equivalent to <existing window name>.
 - b) Let *WDX* be the window structure descriptor identified by *EWN*.
 - c) *WDEF* shall not specify <window partition clause>.
 - d) If *WDX* has a window ordering clause, then *WDEF* shall not specify <window order clause>.
 - e) *WDX* shall not have a window framing clause.
- 11) If *WDEF*'s <window frame clause> specifies <window frame preceding> or <window frame following>, then let *UVS* be the <unsigned value specification> simply contained in the <window frame preceding> or <window frame following>.

Case:

- a) If RANGE is specified, then:
 - i) Case:
 - 1) If *WDEF* contains <window order clause> *WOC*, then let *WDEF**OC* be *WOC*.
 - 2) Otherwise, *WDEF* shall specify an <existing window name> that identifies a window structure descriptor that includes a window ordering clause *WOGC*. Let *WDEF**OC* be *WOGC*.
 - ii) *WDEF**OC* shall contain a single <sort key> *SK*.
 - iii) The declared type of *SK* shall be numeric, datetime, or interval. The declared type of *UVS* shall be numeric if the declared type of *SK* is numeric; otherwise, it shall be an interval type that may be added to or subtracted from the declared type of *SK* according to the Syntax Rules of Subclause 6.30, "<datetime value expression>", and Subclause 6.32, "<interval value expression>", in this part of ISO/IEC 9075.

- b) If ROWS is specified, then the declared type of *UVS* shall be exact numeric with scale 0 (zero).
- 12) The scope of the <new window name> simply contained in *WDEF* consists of any <window definition>s that follow *WDEF* in the <window clause>, together with the <select list> of the <query specification> or <select statement: single row> that simply contains the <window clause>. If the <window clause> is simply contained in a <query specification> that is the <query expression body> of a <query expression> that is a simple table query, then the scope of <new window name> also includes the <order by clause>, if any, of that <query expression>.
- 13) Two window structure descriptors *WD1* and *WD2* are *order-equivalent* if all of the following conditions are met:
- Let $WPCR1_i$, $1 \leq i \leq N1$, and $WPCR2_i$, $1 \leq i \leq N2$, be enumerations of the <window partition column reference>s contained in the window partitioning clauses of *WD1* and *WD2*, respectively, in order from left to right. $N1 = N2$, and, for all i , $WPCR1_i$ and $WPCR2_i$ are equivalent column references.
 - Let $SS1_i$, $1 \leq i \leq M1$, and $SS2_i$, $1 \leq i \leq M2$, be enumerations of the <sort specification>s contained in the window ordering clauses of *WD1* and *WD2*, respectively, in order from left to right. $M1 = M2$, and, for all i , $SS1_i$ and $SS2_i$ contain <sort key>s that are equivalent column references, specify or imply the same <ordering specification>, specify or imply the same <collate clause>, if any, and specify or imply the same <null ordering>.

Access Rules

None.

General Rules

- 1) Let *SL* be the <select list> of the <query specification> or <select statement: single row> that immediately contains *TE*.

Case:

- If *SL* does not simply contain a <window function>, then the <window clause> is disregarded, and the result of *TE* is the result of the last <from clause>, <where clause>, <group by clause> or <having clause> of *TE*.
- Otherwise, let *RTE* be the result of the last <from clause> or <where clause> simply contained in *TE*.

NOTE 178 — Although it is permissible to have a <group by clause> or a <having clause> with a <window clause>, if there are any <window function>s, then the <group by clause> and <having clause> are removed by a syntactic transformation in Subclause 7.12, “<query specification>”, and so are not considered here.

- A window structure descriptor *WDESC* is created for each <window definition> *WDEF*, as follows:
 - WDESC*'s window name is the <new window name> simply contained in *WDEF*.
 - If <existing window name> is specified, then let *EWN* be the <existing window name> simply contained in *WDEF* and let *WDX* be the window structure descriptor identified by *EWN*.

- 3) If <existing window name> is specified and the window ordering clause of *WDX* is present, then the ordering window name of *WDESC* is *EWN*; otherwise, there is no ordering window name.
 - 4) Case:
 - A) If *WDEF* simply contains <window partition clause> *WDEFWPC*, then *WDESC*'s window partitioning clause is *WDEFWPC*.
 - B) If <existing window name> is specified, then *WDESC*'s window partitioning clause is the window partitioning clause of *WDX*.
 - C) Otherwise, *WDESC* has no window partitioning.
 - 5) Case:
 - A) If *WDEF* simply contains <window order clause> *WDEFWOC*, then *WDESC*'s window ordering clause is *WDEFWOC*.
 - B) If <existing window name> is specified, then *WDESC*'s window ordering clause is the window ordering clause of *WDX*.
 - C) Otherwise, *WDESC* has no window ordering.
 - 6) If *WDEF* simply contains <window frame clause> *WDEFWFC*, then *WDESC*'s window framing clause is *WDEFWFC*; otherwise, *WDESC* has no windows framing.
- ii) The result of <window clause> is *RTE*, together with the window structure descriptors defined by the <window clause>.
- 2) Let *WD* be a window structure descriptor.
 - 3) *WD* defines, for each row *R* of *RTE*, the *window partition* of *R* under *WD*, consisting of the collection of rows of *RTE* that are not distinct from *R* in the window partitioning columns of *WD*. If *WD* has no window partitioning clause, then the window partition of *R* is the entire result *RTE*.
 - 4) *WD* also defines the window ordering of the rows of each window partition defined by *WD*, according to the General Rules of [Subclause 10.10](#), “<sort specification list>”, using the <sort specification list> simply contained in *WD*'s window ordering clause. If *WD* has no window ordering clause, then the window ordering is implementation-dependent, and all rows are peers. Although the window ordering of peer rows within a window partition is implementation-dependent, the window ordering shall be the same for all window structure descriptors that are order-equivalent. It shall also be the same for any pair of windows *W1* and *W2* such that *W1* is the ordering window for *W2*.
 - 5) *WD* also defines for each row *R* of *RTE* the window frame *WF* of *R*, consisting of a collection of rows. *WF* is defined as follows.
 Case:
 - a) If *WD* has no window framing clause, then
 Case:
 - i) If the window ordering clause of *WD* is not present, then *WF* is the window partition of *R*.
 - ii) Otherwise, *WF* consists of all rows of the partition of *R* that precede *R* or are peers of *R* in the window ordering of the window partition defined by the window ordering clause.

- b) Otherwise, let *WF* initially be the window partition of *R* defined by *WD*. Let *WFC* be the window framing clause of *WD*. Let *WFB1* be the <window frame bound 1> and let *WFB2* be the <window frame bound 2> contained in *WFC*.

i) If RANGE is specified, then:

- 1) In the following subrules, when performing addition or subtraction to combine a datetime and a year-month interval, if the result would raise the exception condition *data exception — datetime field overflow* because the <primary datetime field> DAY is not valid for the computer value of the <primary datetime field>s YEAR and MONTH, then the <primary datetime field> DAY is set to the last day that is valid for the <primary datetime field>s YEAR and MONTH, and no exception condition is raised.

2) Case:

NOTE 179 — In the following subrules, if *WFB1* specifies UNBOUNDED PRECEDING, then no rows are removed from *WF* by this step. *WFB1* may not be UNBOUNDED FOLLOWING.

- A) If *WFB1* specifies <window frame preceding>, then let *VIP* be the value of the <unsigned value specification>.

Case:

- I) If *VIP* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
- II) Otherwise, let *SK* be the only <sort key> contained in the window ordering clause of *WD*. Let *VSK* be the value of *SK* for the current row.

Case:

- 1) If *VSK* is the null value and if NULLS LAST is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is not the null value.

2) If *VSK* is not the null value, then:

- a) If NULLS FIRST is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is the null value.

b) Case:

- i) If the <ordering specification> contained in the window ordering clause specifies DESC, then let *BOUND* be the value *VSK+VIP*. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is greater than *BOUND*.
- ii) Otherwise, let *BOUND* be the value *VSK-VIP*. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is less than *BOUND*.

- B) If *WFB1* specifies CURRENT ROW, then remove from *WF* all rows that are not peers of the current row and that precede the current row in the window ordering defined by *WD*.

- C) If *WFB1* specifies <window frame following>, then let *VIF* be the value of the <unsigned value specification>.

Case:

- I) If *VIF* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
- II) Otherwise, let *SK* be the only <sort key> contained in the window ordering clause of *WD*. Let *VSK* be the value of *SK* for the current row.

Case:

- 1) If *VSK* is the null value and if *NULLS LAST* is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is not the null value.
- 2) If *VSK* is not the null value, then:
 - a) If *NULLS FIRST* is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is the null value.
 - b) Case:
 - i) If the <ordering specification> contained in the window ordering clause specifies *DESC*, then let *BOUND* be the value *VSK-VIF*. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is greater than *BOUND*.
 - ii) Otherwise, let *BOUND* be the value *VSK+VIF*. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is less than *BOUND*.
- 3) Case:

NOTE 180 — In the following subrules, if *WFB2* specifies *UNBOUNDED FOLLOWING*, then no rows are removed from *WF* by this step. *WFB2* may not be *UNBOUNDED PRECEDING*.

- A) If *WFB2* specifies <window frame preceding>, then let *V2P* be the value of the <unsigned value specification>.

Case:

- I) If *V2P* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
- II) Otherwise, let *SK* be the only <sort key> contained in the window ordering clause of *WD*. Let *VSK* be the value of *SK* for the current row.

Case:

- 1) If *VSK* is the null value and if *NULLS FIRST* is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is not the null value.
- 2) If *VSK* is not the null value, then:
 - a) If *NULLS LAST* is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is the null value.
 - b) Case:

- i) If the <ordering specification> contained in the window ordering clause specifies DESC, then let *BOUND* be the value $VSK+V2P$. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is less than *BOUND*.
 - ii) Otherwise, let *BOUND* be the value $VSK-V2P$. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is greater than *BOUND*.
- B) If *WFB2* specifies CURRENT ROW, then remove from *WF* all rows following the current row in the ordering defined by *WD* that are not peers of the current row.
- C) If *WFB2* specifies <window frame following>, then let *V2F* be the value of the <unsigned value specification>.

Case:

- I) If *V2F* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
- II) Otherwise, let *SK* be the only <sort key> contained in the window ordering clause of *WD*. Let *VSK* be the value of *SK* for the current row.

Case:

- 1) If *VSK* is the null value and if NULLS FIRST is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is not the null value.
- 2) If *VSK* is not the null value, then:
 - a) If NULLS LAST is specified or implied, then remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is the null value.
 - b) Case:
 - i) If the <ordering specification> contained in the <window order clause> specifies DESC, then let *BOUND* be the value $VSK-V2F$. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is less than *BOUND*.
 - ii) Otherwise, let *BOUND* be the value $VSK+V2F$. Remove from *WF* all rows *R2* such that the value of *SK* in row *R2* is greater than *BOUND*.
- ii) If ROWS is specified, then:

1) Case:

NOTE 181 — In the following subrules, if *WFB1* specifies UNBOUNDED PRECEDING, then no rows are removed from *WF* by this step. *WFB1* may not be UNBOUNDED FOLLOWING.

- A) If *WFB1* specifies <window frame preceding>, then let *V1P* be the value of the <unsigned value specification>.

Case:

- I) If *VIP* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
 - II) Otherwise, remove from *WF* all rows that are more than *VIP* rows preceding the current row in the window ordering defined by *WD*.
- B) If *WFB1* specifies **CURRENT ROW**, then remove from *WF* all rows that precede the current row in the window ordering defined by *WD*.

NOTE 182 — This step removes any peers of the current row that precede it in the implementation-dependent window ordering.

- C) If *WFB1* specifies <window frame following>, then let *VIF* be the value of the <unsigned value specification>.

Case:

- I) If *VIF* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
- II) Otherwise, remove from *WF* all rows that precede the current row and all rows that are less than *VIF* rows following the current row in the window ordering defined by *WD*.

NOTE 183 — If *VIF* is zero, then the current row is not removed from *WF* by this step; otherwise, the current row is removed from *WF*.

2) Case:

NOTE 184 — In the following subrules, if *WFB2* specifies **UNBOUNDED FOLLOWING**, then no rows are removed from *WF* by this step. *WFB2* may not be **UNBOUNDED PRECEDING**.

- A) If *WFB2* specifies <window frame preceding>, then let *V2P* be the value of the <unsigned value specification>.

Case:

- I) If *V2P* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.
- II) Otherwise, remove from *WF* all rows that follow the current row and all rows that are less than *V2P* rows preceding the current row in the window ordering defined by *WD*.

NOTE 185 — If *V2P* is zero, then the current row is not removed from *WF* by this step; otherwise, the current row is removed from *WF*.

- B) If *WFB2* specifies **CURRENT ROW**, then remove from *WF* all rows that follow the current row in the window ordering defined by *WD*.

NOTE 186 — This step removes any peers of the current row that follow it in the implementation-dependent window ordering.

- C) If *WFB2* specifies <window frame following>, then let *V2F* be the value of the <unsigned value specification>.

Case:

- I) If *V2F* is negative or the null value, then an exception condition is raised: *data exception — invalid preceding or following size in window function*.

- II) Otherwise, remove from *WF* all rows that are more than *V2F* rows following the current row in the window ordering defined by *WD*.
- iii) If <window frame exclusion> *WFE* is specified, then

Case:

- 1) If EXCLUDE CURRENT ROW is specified and the current row is still a member of *WF*, then remove the current row from *WF*.
- 2) If EXCLUDE GROUP is specified, then remove the current row and any peers of the current row from *WF*.
- 3) If EXCLUDE TIES is specified, then remove any rows other than the current row that are peers of the current row from *WF*.

NOTE 187 — If the current row is already removed from *WF*, then it remains removed from *WF*.

NOTE 188 — If EXCLUDE NO OTHERS is specified, then no additional rows are removed from *WF* by this Rule.

Conformance Rules

- 1) Without Feature T611, “Elementary OLAP operations”, conforming SQL language shall not contain a <window specification>.
- 2) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window clause>.
- 3) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain an <existing window name>.
- 4) Without Feature T301, “Functional dependencies”, in conforming SQL language, if *T* is a grouped table, then each column reference contained in <window clause> that references a column of *T* shall be a reference to a grouping column of *T* or be contained in an aggregated argument of a <set function specification>.
- 5) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window frame exclusion>.

NOTE 189 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

7.12 <query specification>

This Subclause is modified by Subclause 7.1, “<query specification>”, in ISO/IEC 9075-4.

Function

Specify a table derived from the result of a <table expression>.

Format

```
<query specification> ::=
    SELECT [ <set quantifier> ] <select list> <table expression>

<select list> ::=
    <asterisk>
    | <select sublist> [ { <comma> <select sublist> }... ]

<select sublist> ::=
    <derived column>
    | <qualified asterisk>

<qualified asterisk> ::=
    <asterisked identifier chain> <period> <asterisk>
    | <all fields reference>

<asterisked identifier chain> ::=
    <asterisked identifier> [ { <period> <asterisked identifier> }... ]

<asterisked identifier> ::=
    <identifier>

<derived column> ::=
    <value expression> [ <as clause> ]

<as clause> ::=
    [ AS ] <column name>

<all fields reference> ::=
    <value expression primary> <period> <asterisk>
    [ AS <left paren> <all fields column name list> <right paren> ]

<all fields column name list> ::=
    <column name list>
```

Syntax Rules

- 1) Let *T* be the result of the <table expression>.
- 2) Let *TQS* be the table that is the result of a <query specification>.
- 3) Case:

- a) If the <select list> “*” is simply contained in a <table subquery> that is immediately contained in an <exists predicate>, then the <select list> is equivalent to a <value expression> that is an arbitrary <literal>.
 - b) Otherwise, the <select list> “*” is equivalent to a <value expression> sequence in which each <value expression> is a column reference that references a column of *T* and each column of *T* is referenced exactly once. The columns are referenced in the ascending sequence of their ordinal position within *T*.
- 4) The degree of the table specified by a <query specification> is equal to the cardinality of the <select list>.
 - 5) If a <set quantifier> DISTINCT is specified, then each column of *T* is an operand of a grouping operation. The Syntax Rules of [Subclause 9.10, “Grouping operations”](#), apply.
 - 6) The ambiguous case of an <all fields reference> whose <value expression primary> takes the form of an <asterisked identifier chain> shall be analyzed first as an <asterisked identifier chain> to resolve the ambiguity.
 - 7) If <asterisked identifier chain> is specified, then:
 - a) Let *IC* be an <asterisked identifier chain>.
 - b) Let *N* be the number of <asterisked identifier>s immediately contained in *IC*.
 - c) Let *I_i*, $1 \text{ (one)} \leq i \leq N$, be the <asterisked identifier>s immediately contained in *IC*, in order from left to right.
 - d) Let *PIC₁* be *I₁*. For each *J* between 2 and *N*, let *PIC_J* be *PIC_{J-1}.I_J*. *PIC_J* is called the *J*-th *partial identifier chain* of *IC*.
 - e) Let *M* be the minimum of *N* and 3.
 - f) For at most one *J* between 1 and *M*, *PIC_J* is called the *basis* of *IC*, and *J* is called the *basis length* of *IC*. The *referent* of the basis is a table *T*, a column *C* of a table, or an SQL parameter *SP*. The *basis* and *basis scope* of *IC* are defined in terms of a *candidate basis*, according to the following rules:
 - i) 04 If *IC* is contained in the scope of a <routine name> whose associated <SQL parameter declaration list> includes an SQL parameter *SP* whose <SQL parameter name> is equivalent to *I₁*, then *PIC₁* is a candidate basis of *IC*, and the scope of *PIC₁* is the scope of *SP*.
 - ii) If *N* = 2 and *PIC₁* is equivalent to the <qualified identifier> of a <routine name> *RN* whose scope contains *IC* and whose associated <SQL parameter declaration list> includes an SQL parameter *SP* whose <SQL parameter name> is equivalent to *I₂*, then *PIC₂* is a candidate basis of *IC*, the scope of *PIC₂* is the scope of *SP*, and the referent of *PIC₂* is *SP*.
 - iii) If *N* > 2 and *PIC₁* is equivalent to the <qualified identifier> of a <routine name> *RN* whose scope contains *IC* and whose associated <SQL parameter declaration list> includes a refinable SQL parameter *SP* whose <SQL parameter name> is equivalent to *I₂*, then *PIC₂* is a candidate basis of *IC*, the scope of *PIC₂* is the scope of *SP*, and the referent of *PIC₂* is *SP*.
 - iv) If *N* = 2 and *PIC₁* is equivalent to an exposed <correlation name> that is in scope, then let *EN* be the exposed <correlation name> that is equivalent to *PIC₁* and has innermost scope. If the table associated with *EN* has a column *C* of row type whose <identifier> is equivalent to *I₂*, then *PIC₂* is a candidate basis of *IC* and the scope of *PIC₂* is the scope of *EN*.

- v) If $N > 2$ and PIC_1 is equivalent to an exposed <correlation name> that is in scope, then let EN be the exposed <correlation name> that is equivalent to PIC_1 and has innermost scope. If the table associated with EN has a column C of row type or structured type whose <identifier> is equivalent to I_2 , then PIC_2 is a candidate basis of IC and the scope of PIC_2 is the scope of EN .
 - vi) If $N \leq 3$ and PIC_N is equivalent to an exposed <table or query name> that is in scope, then let EN be the exposed <table or query name> that is equivalent to PIC_N and has the innermost scope. PIC_N is a candidate basis of IC , and the scope of PIC_N is the scope of EN .
 - vii) There shall be exactly one candidate basis CB with innermost scope. The basis of IC is CB . The basis scope is the scope of CB .
- g) Case:
- i) If the basis is a <table or query name> or <correlation name>, then let TQ be the table associated with the basis. The <select sublist> is equivalent to a <value expression> sequence in which each <value expression> is a column reference CR that references a column of TQ that is not a common column of a <joined table>. Each column of TQ that is not a referenced common column shall be referenced exactly once. The columns shall be referenced in the ascending sequence of their ordinal positions within TQ .
 - ii) Otherwise let BL be the length of the basis of IC .

Case:

- 1) If $BL = N$, then the <select sublist> $IC . *$ is equivalent to $(IC) . *$.
- 2) Otherwise, the <select sublist> $IC . *$ is equivalent to:

$$(PIC_{BL}) . I_{BL+1} . \dots . I_N . *$$

NOTE 190 — The equivalent syntax in either case will be analyzed as <all fields reference> ::= <value expression primary> <period> <asterisk>

- 8) The data type of the <value expression primary> VEP specified in an <all fields reference> AFR shall be some row type VER . Let n be the degree of VER . Let F_1, \dots, F_N be the field names of VER .

Case:

- a) If <all fields column name list> $AFCNL$ is specified, then the number of <column name>s simply contained in $AFCNL$ shall be n . Let $AFCN_i$, $1 \text{ (one)} \leq i \leq n$, be these <column name>s in order from left to right. AFR is equivalent to

```

VEP . F1
AS AFCN1, ... , VEP
. Fn AS
AFCNn)

```

- b) Otherwise, AFR is equivalent to:

```

VEP . F1 , ... ,
VEP . Fn

```

- 9) Let C be some column. Let QS be the <query specification>. Let DC_i , for i ranging from 1 (one) to the number of <derived column>s inclusively, be the i -th <derived column> simply contained in the <select list> of QS . For all i , C is an underlying column of DC_i , and of any column reference that identifies DC_i , if and only if C is an underlying column of the <value expression> of DC_i , or C is an underlying column of the <table expression> immediately contained in QS .
- 10) Each column reference contained in a <window function> shall unambiguously reference a column of T .
- 11) If both of the following two conditions are satisfied, then QS is a *grouped, windowed query*:
 - a) T is a grouped table.
 - b) Some <derived column> simply contained in QS simply contains a <window function>.
- 12) A grouped, windowed query GWQ is transformed to an equivalent <query specification> as follows:
 - a) If GWQ contains an <in-line window specification>, then apply the syntactic transformation specified in Subclause 6.10, "<window function>".
 - b) If the <select list> of GWQ immediately contains <asterisk> or simply contains <qualified asterisk>, then apply the syntactic transformations specified in Subclause 7.12, "<query specification>".
 - c) Let $GWQ2$ be the result of the preceding transformations, if any.
 - d) Let SL , FC , WC , GBC , HC , and WIC be the <select list>, <from clause>, <where clause>, <group by clause>, <having clause>, and <window clause>, respectively, of $GWQ2$. If any of <where clause>, <group by clause>, or <having clause>, are missing, then let WC , GBC , and HC , respectively, be a zero-length string. Let SQ be the <set quantifier> immediately contained in the <query specification> of $GWQ2$, if any; otherwise, let SQ be a zero-length string.

NOTE 191 — $GWQ2$ can not lack a <window clause>, since the syntactic transformation of Subclause 6.10, "<window function>", will create one if there is not one in GWQ already.
 - e) Let NI be the number of <set function specification>s simply contained in $GWQ2$.
 - f) Let SFS_i , $1 \text{ (one)} \leq i \leq NI$, be an enumeration of the <set function specification>s simply contained in $GWQ2$.
 - g) Let $SFSI_i$, $1 \text{ (one)} \leq i \leq NI$, be a list of <identifier>s that are distinct from each other and distinct from all <identifier>s contained in $GWQ2$.
 - h) If $NI = 0$ (zero), then let $SFSL$ be a zero-length string; otherwise, let $SFSL$ be:

$$\begin{aligned} SFS_1 & \text{ AS } SFSI_1, \\ SFS_2 & \text{ AS } SFSI_2, \dots, \\ SFS_{NI} & \text{ AS } SFSI_{NI} \end{aligned}$$
 - i) Let $HCNEW$ be obtained from HC by replacing each <set function specification> SFS_i by the corresponding <identifier> $SFSI_i$.
 - j) Let $N2$ be the number of <column reference>s that are contained in SL or WIC without an intervening <query expression> or <set function specification>.
 - k) Let CR_j , $1 \text{ (one)} \leq j \leq N2$, be an enumeration of the <column reference>s that are contained in SL or WIC without an intervening <query expression> or <set function specification>.

- l) Let CRI_j , $1 \text{ (one)} \leq j \leq N2$, be a list of <identifier>s that are distinct from each other, distinct from all identifiers in $GWQ2$, and distinct from all $SFSI_i$.
- m) If $N2 = 0 \text{ (zero)}$, then let CRL be a zero-length string; otherwise, let CRL be:

$CR_1 \text{ AS } CRI_1, CR_2 \text{ AS } CRI_2, \dots, CR_{N2} \text{ AS } CRI_{N2}$

- n) Let $N3$ be the number of <derived column>s simply contained in SL that do not specify <as clause>.
- o) Let $DCOL_k$, $1 \text{ (one)} \leq k \leq N3$, be the <derived column>s simply contained in SL that do not specify an <as clause>. For each k , let $COLN_k$ be the <column name> determined as follows.

Case:

- i) If $DCOL_k$ is a single column reference, then $COLN_k$ is the <column name> of the column designated by the column reference.
- ii) 04 If $DCOL_k$ is a single SQL parameter reference, then $COLN_k$ is the <SQL parameter name> of the SQL parameter designated by the SQL parameter reference.
- iii) Otherwise, $COLN_k$ is an implementation-dependent <column name>.
- p) Let $SL2$ be obtained from SL by replacing each <derived column> $DCOL_k$ by

$DCOL_k \text{ AS } COLN_k$

- q) Let $GWQN$ be an arbitrary <identifier>.
- r) Let $SLNEW$ be the <select list> obtained from $SL2$ by replacing each simply contained <set function specification> SFS_i by $GWQN.SFSI_i$ and replacing each <column reference> CR_j that is contained without an intervening <query expression> or <set function specification> by $GWQN.CRI_j$.
- s) Let $WICNEW$ be the <window clause> obtained from WIC by replacing each <set function specification> SFS_i by $GWQN.SFSI_i$ and by replacing each <column reference> CR_j by $GWQN.CRI_j$.
- t) If either $SFSL$ or CRL is a zero-length string, then let $COMMA$ be a zero-length string; otherwise, let $COMMA$ be “,” (a <comma>).
- u) GWQ is equivalent to the following <query specification>:

```
SELECT SLNEW
FROM ( SELECT SQ SFSL COMMA CRL
      FC
      WC
      GBC
      HC ) AS GWQN
WICNEW
```

- 13) A <query specification> is *possibly non-deterministic* if any of the following conditions are true:

- a) The <set quantifier> DISTINCT is specified and one of the columns of T has a data type of character string, user-defined type, TIME WITH TIME ZONE, or TIMESTAMP WITH TIME ZONE.
- b) The <query specification> generally contains a <value expression>, <query specification>, or <query expression> that is possibly non-deterministic.

- c) The <select list>, <having clause>, or <window clause> contains a reference to a column C of T that has a data type of character string, user-defined type, TIME WITH TIME ZONE, or TIMESTAMP WITH TIME ZONE, and the functional dependency $G \twoheadrightarrow C$, where G is the set consisting of the grouping columns of T , holds in T .
- 14) If <table expression> does not immediately contain a <group by clause> and <table expression> is simply contained in a <query specification> that is the aggregation query of some <set function specification>, then GROUP BY () is implicit.

NOTE 192 — “aggregation query” is defined in Subclause 6.9, “<set function specification>”.

- 15) If T is a grouped table, then let G be the set of grouping columns of T . In each <value expression> contained in <select list>, each column reference that references a column of T shall reference some column C that is functionally dependent on G or shall be contained in an aggregated argument of a <set function specification> whose aggregation query is QS .

NOTE 193 — See also the Syntax Rules of Subclause 6.7, “<column reference>”.

- 16) Each column of TQS has a column descriptor that includes a data type descriptor that is the same as the data type descriptor of the <value expression> simply contained in the <derived column> defining that column.
- 17) Case:
- If the i -th <derived column> in the <select list> specifies an <as clause> that contains a <column name> CN , then the <column name> of the i -th column of the result is CN .
 - If the i -th <derived column> in the <select list> does not specify an <as clause> and the <value expression> of that <derived column> is a single column reference, then the <column name> of the i -th column of the result is the <column name> of the column designated by the column reference.
 - 04 If the i -th <derived column> in the <select list> does not specify an <as clause> and the <value expression> of that <derived column> is a single SQL parameter reference, then the <column name> of the i -th column of the result is the <SQL parameter name> of the SQL parameter designated by the SQL parameter reference.
 - Otherwise, the <column name> of the i -th column of the <query specification> is implementation-dependent.
- 18) A column C of TQS is *readily known not null* if C is defined (after performing syntactic transformations to eliminate <asterisk> and <qualified asterisk>), by a <derived column> that immediately contains a <value expression> that is a <column reference> CR such that the column referenced by CR is a readily known not null column of the qualifying table of CR .
- 19) A column $CKNN$ of TQS is *known not null* if and only if at least one of the following conditions applies:
- $CKNN$ is readily known not null.
 - The SQL-implementation supports Feature T101, “Enhanced nullability determination” and $CKNN$ is not defined by a <derived column> containing any of the following:
 - A column reference for a column C that is possibly nullable.
 - An <indicator parameter>.
 - An <indicator variable>.
 - 04 A <dynamic parameter specification>.

- v) An SQL parameter.
 - vi) A <routine invocation>, <method reference>, or <method invocation> whose subject routine is an SQL-invoked routine that either is an SQL routine or is an external routine that specifies or implies PARAMETER STYLE SQL.
 - vii) A <query expression>.
 - viii) CAST (NULL AS X) (where X represents a <data type> or a <domain name>).
 - ix) A <window function> whose <window function type> does not contain <rank function type>, ROW_NUMBER, or an <aggregate function> that simply contains COUNT.
 - x) CURRENT_USER, CURRENT_ROLE, or SYSTEM_USER.
 - xi) A <set function specification> that does not simply contain COUNT.
 - xii) A <case expression>.
 - xiii) A <field reference>.
 - xiv) An <array element reference>.
 - xv) A <multiset element reference>.
 - xvi) A <dereference operation>.
 - xvii) A <reference resolution>.
 - xviii) A <comparison predicate>, <between predicate>, <in predicate>, or <quantified comparison predicate> *P* such that the declared type of a field of a <row value predicand> that is simply contained in *P* is a row type, a user-defined type, an array type, or a multiset type.
 - xix) A <member predicate>.
 - xx) A <submultiset predicate>.
- c) An implementation-defined rule by which the SQL-implementation can correctly deduce that the value of the column cannot be null.

20) Let *TREF* be the <table reference>s that are simply contained in the <from clause> of the <table expression>. The *simply underlying tables* of the <query specification> are the <table or query name>s and <derived table>s contained in *TREF* without an intervening <derived table>.

21) The terms *key-preserving* and *one-to-one* are defined as follows:

- a) Let *UT* denote some simply underlying table of *QS*, let *UTCOLS* be the set of columns of *UT*, let *QSCOLS* be the set of columns of *QS*, and let *QSCN* be an exposed range variable for *UT* whose scope clause is *QS*.
- b) *QS* is said to be *key-preserving with respect to UT* if there is some strong candidate key *CKUT* of *UT* such that every member of *CKUT* has some counterpart under *QSCN* in *QSCOLS*.

NOTE 194 — “strong candidate key” is defined in Subclause 4.19, “Candidate keys”.

NOTE 195 — “Counterpart” is defined in Subclause 4.18.2, “General rules and definitions”. It follows from this condition that every row in *QS* corresponds to exactly one row in *UT*, namely that row in *UT* that has the same combined value in the columns of *CKUT* as the row in *QS*. There may be more than one row in *QS* that corresponds to a single row in *UT*.

- c) *QS* is said to be *one-to-one with respect to UT* if and only if *QS* is key-preserving with respect to *UT*, *UT* is updatable, and there is some strong candidate key *CKQS* of *QS* such that every member of *CKQS* is a counterpart under *UT* of some member of *UTCOLS*.

NOTE 196 — It follows from this condition that every row in *UT* corresponds to at most one row in *QS*, namely that row in *QS* that has the same combined value in the columns of *CKQS* as the row in *UT*.

22) A <query specification> is *potentially updatable* if and only if the following conditions hold:

- a) DISTINCT is not specified.
- b) Of those <derived column>s in the <select list> that are column references that have a counterpart in a base table, no column of a table is referenced more than once in the <select list>.
- c) The <table expression> immediately contained in *QS* does not simply contain an explicit or implicit <group by clause> or a <having clause>.

23) If a <query specification> *QS* is potentially updatable, then

Case:

- a) If the <from clause> of the <table expression> specifies exactly one <table reference>, then a column of *QS* is said to be a *potentially updatable column* if it has a counterpart in *TR* that is updatable.

NOTE 197 — The notion of updatable columns of table references is defined in Subclause 7.6, “<table reference>”.

- b) Otherwise, a column of *QS* is said to be a *potentially updatable column* if it has a counterpart in some updatable column of some simply underlying table *UT* of *QS* such that *QS* is one-to-one with respect to *UT*.

24) A <query specification> is *updatable* if it is potentially updatable and it has at least one potentially updatable column.

25) A <query specification> *QS* is *simply updatable* if the following conditions hold:

- a) *QS* is updatable.
- b) The <from clause> immediately contained in the <table expression> immediately contained in *QS* contains exactly one <table reference>, and the table referenced by that <table reference> is simply updatable.
- c) Every result column of *QS* is updatable.
- d) If the <table expression> immediately contained in *QS* immediately contains a <where clause> *WC*, then no leaf generally underlying table of *QS* is a generally underlying table of any <query expression> contained in *WC*.

26) A <query specification> *QS* is *insertable-into* if and only if every simply underlying table of *QS* is insertable-into.

27) A column *C* of *QS* is updatable if at least one of the following is true:

- a) *QS* is simply updatable.
- b) *QS* is updatable, *C* is potentially updatable, and the SQL implementation supports Feature T111, “Updatable joins, unions, and columns”.

28) The row type *RT* of *TQS* is defined by the sequence of (<field name>, <data type>) pairs indicated by the sequence of column descriptors of *TQS* taken in order.

Access Rules

None.

General Rules

1) Case:

- a) If *T* is not a grouped table, then each <value expression> is applied to each row of *T* yielding a table *TEMP* of *M* rows, where *M* is the cardinality of *T*. The *i*-th column of the table contains the values derived by the evaluation of the *i*-th <value expression>.
- b) If *T* is a grouped table, then
Case:
 - i) If *T* has 0 (zero) groups, then let *TEMP* be an empty table.
 - ii) If *T* has one or more groups, then each <value expression> is applied to each group of *T* yielding a table *TEMP* of *M* rows, where *M* is the number of groups in *T*. The *i*-th column of *TEMP* contains the values derived by the evaluation of the *i*-th <value expression>. When a <value expression> is applied to a given group of *T*, that group is the argument source of each <set function specification> whose aggregation query is *QS*.

2) Case:

- a) If the <set quantifier> DISTINCT is not specified, then the result of the <query specification> is *TEMP*.
- b) If the <set quantifier> DISTINCT is specified, then the result of the <query specification> is the table derived from *TEMP* by the elimination of all redundant duplicate rows. If the most specific type of any column is character string, datetime with time zone, or a user-defined type, then the precise values in those columns are chosen in an implementation-dependent fashion.

Conformance Rules

- 1) Without Feature F801, “Full set function”, conforming SQL language shall not contain a <query specification> *QS* that contains more than one <set quantifier> containing DISTINCT, unless such <set quantifier> is contained in a <query expression> contained in *QS*.
- 2) Without Feature T051, “Row types”, conforming SQL language shall not contain an <all fields reference>.
- 3) Without Feature T301, “Functional dependencies”, in conforming SQL language, if *T* is a grouped table, then in each <value expression> contained in the <select list>, each <column reference> that references a column of *T* shall reference a grouping column or be specified in an aggregated argument of a <set function specification>.
- 4) Without Feature T325, “Qualified SQL parameter references”, conforming SQL language shall not contain an <asterisked identifier chain> whose referent is an SQL parameter and whose first <identifier> is the <qualified identifier> of a <routine name>.
- 5) Without Feature T053, “Explicit aliases for all-fields reference”, conforming SQL language shall not contain an <all fields column name list>.

NOTE 198 — If a <set quantifier> DISTINCT is specified, then the Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

- 6) Without Feature T285, “Enhanced derived column names”, in conforming SQL language, if any <derived column> in a <select list> does not specify an <as clause> and the <value expression> of that <derived column> is not a single column reference, then the <column name> of that column is implementation-dependent.

7.13 <query expression>

This Subclause is modified by Subclause 7.2, “<query expression>”, in ISO/IEC 9075-14.

Function

Specify a table.

Format

```

<query expression> ::=
  [ <with clause> ] <query expression body>
    [ <order by clause> ] [ <fetch first clause> ]

14 <with clause> ::=
  WITH [ RECURSIVE ] <with list>

<with list> ::=
  <with list element> [ { <comma> <with list element> }... ]

<with list element> ::=
  <query name> [ <left paren> <with column list> <right paren> ]
    AS <table subquery> [ <search or cycle clause> ]

<with column list> ::=
  <column name list>

<query expression body> ::=
  <query term>
  | <query expression body> UNION [ ALL | DISTINCT ]
    [ <corresponding spec> ] <query term>
  | <query expression body> EXCEPT [ ALL | DISTINCT ]
    [ <corresponding spec> ] <query term>

<query term> ::=
  <query primary>
  | <query term> INTERSECT [ ALL | DISTINCT ]
    [ <corresponding spec> ] <query primary>

<query primary> ::=
  <simple table>
  | <left paren> <query expression body>
    [ <order by clause> ] [ <fetch first clause> ] <right paren>

<simple table> ::=
  <query specification>
  | <table value constructor>
  | <explicit table>

<explicit table> ::=
  TABLE <table or query name>

<corresponding spec> ::=
  CORRESPONDING [ BY <left paren> <corresponding column list> <right paren> ]

<corresponding column list> ::=

```

```
<column name list>

<order by clause> ::=
  ORDER BY <sort specification list>

<fetch first clause> ::=
  FETCH FIRST [ <unsigned integer> ] { ROW | ROWS } ONLY
```

Syntax Rules

- 1) Let *QE* be the <query expression>.
- 2) If <with clause> is specified, then:
 - a) If a <with clause> *WC* immediately contains RECURSIVE, then *WC*, its <with list>, and its <with list element>s are said to be *potentially recursive*. Otherwise they are said to be *non-recursive*.
 - b) Let *n* be the number of <with list element>s. For each *i*, $1 \text{ (one)} \leq i < n$, for each *j*, $i < j \leq n$, the *j*-th <with list element> shall not immediately contain a <query name> that is equivalent to the <query name> immediately contained in the *i*-th <with list element>.
 - c) If the <with clause> is non-recursive, then for all *i* between 1 (one) and *n*, the scope of the <query name> *WQN* immediately contained in the *i*-th <with list element> *WLE_i* is the <query expression> simply contained in every <with list element> *WLE_k*, where *k* ranges from *i*+1 to *n*, and the <query expression body> immediately contained in <query expression>. A <table or query name> contained in this scope that immediately contains *WQN* is a *query name in scope*.
 - d) If the <with clause> is potentially recursive, then for all *i* between 1 (one) and *n*, the scope of the <query name> *WQN* immediately contained in the *i*-th <with list element> *WLE_i* is the <query expression> simply contained in every <with list element> *WLE_k*, where *k* ranges from 1 (one) to *n*, and the <query expression body> immediately contained in <query expression>. A <table or query name> contained in this scope that immediately contains *WQN* is a *query name in scope*.
 - e) For every <with list element> *WLE*, let *WQE* be the <query expression> specified by *WLE* and let *WQT* be the table defined by *WQE*.
 - i) If any two columns of *WQT* have equivalent names or if *WLE* is potentially recursive, then *WLE* shall specify a <with column list>. If *WLE* specifies a <with column list> *WCL*, then:
 - 1) Equivalent <column name>s shall not be specified more than once in *WCL*.
 - 2) The number of <column name>s in *WCL* shall be the same as the degree of *WQT*.
 - ii) Every column of a character string type in *WQT* shall have a declared type collation.
 - f) A *query name dependency graph QNDG* of a potentially recursive <with list> *WL* is a directed graph such that, for *i* ranging from 1 (one) to the number of <query name>s simply contained in *WL*:
 - i) Each node represents a <query name> *WQN_i* immediately contained in a <with list element> *WLE_i* of *WL*.
 - ii) Each arc from a node *WQN_i* to a node *WQN_j* represents the fact that *WQN_j* is referenced by a <query name> contained in the <query expression> simply contained in *WLE_i*. *WQN_i* is said to *depend immediately* on *WQN_j*.

- g) For a potentially recursive <with list> WL with n elements, and for i ranging from 1 (one) to n , let WLE_i be the i -th <with list element> of WL , let WQN_i be the <query name> immediately contained in WLE_i , let WQE_i be the <query expression> simply contained in WLE_i , let WQT_i be the table defined by WQE_i , and let $QNDG$ be the query name dependency graph of WL .

- i) WL is said to be *recursive* if $QNDG$ contains at least one cycle.

Case:

- 1) If $QNDG$ contains an arc from WQN_i to itself, then WLE_i , WQN_i , and WQT_i are said to be *recursive*. WQN_i is said to belong to the *stratum* of WQE_i .
- 2) If $QNDG$ contains a cycle comprising WQN_i, \dots, WQN_k , with $k \neq i$, then it is said that WQN_i, \dots, WQN_k are *recursive* and *mutually recursive* to each other, WQT_i, \dots, WQT_k are *recursive* and *mutually recursive* to each other, and WLE_i, \dots, WLE_k are *recursive* and *mutually recursive* to each other.

For each j ranging from i to k , WQN_j belongs to the stratum of WQE_i, \dots , and WQE_k .

- 3) Among the WQE_i, \dots, WQE_k of a given stratum, there shall be at least one <query expression>, say WQE_j , such that:
 - A) WQE_j is a <query expression body> that immediately contains UNION.
 - B) WQE_j has one operand that does not contain a <query name> referencing any of WQN_i, \dots, WQN_k . This operand is said to be the *non-recursive operand* of WQE_j .
 - C) WQE_j is said to be an *anchor expression*, and WQN_j an *anchor name*.
 - D) Let $CCCG$ be the subgraph of $QNDG$ that contains no nodes other than WQN_i, \dots, WQN_k . For any anchor name WQN_j , remove the arcs to those query names WQN_l that are referenced by any <query name> contained in WQE_j . The remaining graph $SCCGP$ shall not contain a cycle.

- ii) If WLE_i is recursive, then

Case:

- 1) If WQE_i contains at most one WQN_k that belongs to the stratum of WQE_i , then WLE_i is *linearly recursive*.
- 2) Otherwise, let WQE_i contain any two <query name>s referencing WQN_k and WQN_l , both of which belong to the stratum of WQE_i .

Case:

- A) WLE_i is *linearly recursive* if each of the following conditions is satisfied:

- I) WQE_i does not contain a <table reference list> that contains <query name>s referencing both WQN_k and WQN_l .
- II) WQE_i does not contain a <joined table> of which $TR1$ and $TR2$ are the first and second <table reference>s, or the <table reference> and <table factor>, respec-

tively, and such that $TR1$ and $TR2$ contain <query name>s referencing WQN_k and WQN_l , respectively.

- III) WQE_i does not contain a <table expression> that immediately contains a <from clause> that contains WQN_k , and immediately contains a <where clause> containing a <query expression> that contains a <query name> referencing WQN_l .

B) Otherwise, WLE_i is said to be *non-linearly recursive*.

- iii) For each WLE_i , for i ranging from 1 (one) to n , and for each WQN_j that belongs to the stratum of WQE_i :
 - 1) WQE_i shall not contain a <query expression body> that contains a <query name> referencing WQN_j and immediately contains EXCEPT where the right operand of EXCEPT contains WQN_j .
 - 2) WQE_i shall not contain a <routine invocation> with an <SQL argument list> that contains one or more <SQL argument>s that immediately contain a <value expression> that contains a <query name> referencing WQN_j .
 - 3) WQE_i shall not contain a <query expression> TSQ that contains a <query name> referencing WQN_j , unless TSQ is simply contained in a <derived table> that is immediately contained in a <table primary> that is immediately contained in a <table reference> that is immediately contained in a <from clause> that is immediately contained in a <table expression> that is immediately contained in a <query specification> that constitutes a <simple table> that constitutes a <query primary> that constitutes a <query term> that is immediately contained in a <query expression body> that is WQE_i .
 - 4) WQE_i shall not contain a <query specification> QS such that:
 - A) QS immediately contains a <table expression> TE that contains a <query name> referencing WQN_j , and
 - B) QS immediately contains a <select list> SL or TE immediately contains a <having clause> HC and SL or TE contain a <set function specification>.
 - 5) WQE_i shall not contain a <query expression body> that contains a <query name> referencing WQN_j and simply contains INTERSECT ALL or EXCEPT ALL.
 - 6) WQE_i shall not contain a <qualified join> QJ in which:
 - A) QJ immediately contains a <join type> that specifies FULL and a <table reference> or <table factor> that contains a <query name> referencing WQN_j .
 - B) QJ immediately contains a <join type> that specifies LEFT and a <table factor> following the <join type> that contains a <query name> referencing WQN_j .
 - C) QJ immediately contains a <join type> that specifies RIGHT and a <table reference> preceding the <join type> that contains a <query name> referencing WQN_j .
 - 7) WQE_i shall not contain a <natural join> QJ in which:

- A) *QJ* immediately contains a <join type> that specifies FULL and a <table reference> or <table primary> that contains a <query name> referencing *WQN_j*.
 - B) *QJ* immediately contains a <join type> that specifies LEFT and a <table primary> following the <join type> that contains a <query name> referencing *WQN_j*.
 - C) *QJ* immediately contains a <join type> that specifies RIGHT and a <table reference> preceding the <join type> that contains a <query name> referencing *WQN_j*.
- iv) If *WLE_i* is recursive, then *WLE_i* shall be linearly recursive.
- v) *WLE_i* is said to be *expandable* if all of the following are true:
- 1) *WLE_i* is recursive.
 - 2) *WLE_i* is linearly recursive.
 - 3) *WQE_i* is a <query expression body> that immediately contains UNION or UNION ALL. Let *WQEB_i* be the <query expression body> immediately contained in *WQE_i*. Let *QEL_i* and *QTR_i* be the <query expression body> and the <query term> immediately contained in *WQEB_i*. *WQN_i* shall not be contained in *QEL_i*, and *QTR_i* shall be a <query specification>.
 - 4) *WQN_i* is not mutually recursive.
- h) If a <with list element> *WLE* is not expandable, then it shall not immediately contain a <search or cycle clause>.
- 3) Let *T* be the table specified by the <query expression>.
- 4) The <explicit table>
- ```
TABLE <table or query name>
```
- is equivalent to the <table subquery>
- ```
( SELECT * FROM <table or query name> )
```
- 5) Let *set operator* be UNION ALL, UNION DISTINCT, EXCEPT ALL, EXCEPT DISTINCT, INTERSECT ALL, or INTERSECT DISTINCT.
- 6) If UNION, EXCEPT, or INTERSECT is specified and neither ALL nor DISTINCT is specified, then DISTINCT is implicit.
- 7) A <query expression> *QE1* is *simply updatable* if it does not specify a <fetch first clause> and, for every <query expression> or <query specification> *QE2* that is simply contained in the <query expression body> of *QE1*, all of the following are true:
- a) *QE1* contains *QE2* without an intervening <query expression body> that specifies UNION ALL, UNION DISTINCT, EXCEPT ALL, or EXCEPT DISTINCT.
 - b) *QE1* contains *QE2* without an intervening <query term> that specifies INTERSECT.
 - c) *QE2* is simply updatable.
- 8) A <query expression> *QE1* is *updatable* if it does not specify a <fetch first clause> and, for every <simple table> *QE2* that is simply contained in *QE1*, all of the following are true:

- a) *QE2* is not a <table value constructor>.
 - b) *QE1* contains *QE2* without an intervening <query expression body> that specifies UNION DISTINCT, EXCEPT ALL, or EXCEPT DISTINCT.
 - c) If *QE1* simply contains a <query expression body> *QEB* that specifies UNION ALL, then:
 - i) *QEB* immediately contains a <query expression body> *LO* and a <query term> *RO* such that no leaf generally underlying table of *LO* is also a leaf generally underlying table of *RO*.
 - ii) For every column of *QEB*, the underlying columns in the tables identified by *LO* and *RO*, respectively, are either both updatable or not updatable.
 - d) *QE1* contains *QE2* without an intervening <query term> that specifies INTERSECT.
 - e) *QE2* is updatable.
- 9) A table specified by a <query name> immediately contained in a <with list element> *WLE* is *updatable* if and only if the <query expression> simply contained in *WLE* is updatable.
- 10) A table specified by a <query name> immediately contained in a <with list element> *WLE* is *simply updatable* if and only if the <query expression> simply contained in *WLE* is simply updatable.
- 11) A <query expression> *QE1* is *insertable-into* if it does not specify a <fetch first clause> and the <query expression body> of *QE1* is a <query primary> that is one of the following:
- a) An insertable-into <query specification>.
 - b) An <explicit table> that identifies a table that is insertable-into.
 - c) Of the form <left paren> <query expression body> <right paren>, where the parenthesized <query expression body> recursively satisfies this condition.
- 12) A table specified by a <query name> immediately contained in a <with list element> *WLE* is *insertable-into* if the <query expression> simply contained in *WLE* is insertable-into.
- 13) For every <simple table> *ST* contained in *QE*,
- Case:
- a) If *ST* is a <query specification> *QS*, then the column descriptor of each column of *ST* is the same as the column descriptor of the corresponding column of *QS*.
 - b) If *ST* is an <explicit table> *ET*, then the column descriptor of each column of *ST* is the same as the column descriptor of the corresponding column of the table identified by the <table or query name> contained in *ET*.
 - c) Otherwise, the column descriptor of each column of *ST* is the same as the column descriptor of the corresponding column of the <table value constructor> immediately contained in *ST*.
- 14) For every <query primary> *QP* contained in *QE*,
- Case:
- a) If *QP* is a <simple table> *ST*, then the column descriptor of each column of *QP* is the same as the column descriptor of the corresponding column of *ST*.
 - b) Otherwise, the column descriptor of each column of *QP* is the same as the column descriptor of the corresponding column of the <query expression body> immediately contained in *QP*.

- 15) If a set operator is specified in a <query term> or a <query expression body>, then:
- Let $T1$, $T2$, and TR be respectively the first operand, the second operand, and the result of the <query term> or <query expression body>.
 - Let $TN1$ and $TN2$ be the effective names for $T1$ and $T2$, respectively.
 - If the set operator is UNION DISTINCT, EXCEPT ALL, EXCEPT DISTINCT, INTERSECT ALL, or INTERSECT DISTINCT, then each column of $T1$ and $T2$ is an operand of a grouping operation. The Syntax Rules of [Subclause 9.10, “Grouping operations”](#), apply.
- 16) If a set operator is specified in a <query term> or a <query expression body>, then let OP be the set operator.

Case:

- If CORRESPONDING is specified, then:
 - Within the columns of $T1$, equivalent <column name>s shall not be specified more than once and within the columns of $T2$, equivalent <column name>s shall not be specified more than once.
 - At least one column of $T1$ shall have a <column name> that is the <column name> of some column of $T2$.
 - Case:
 - If <corresponding column list> is not specified, then let SL be a <select list> of those <column name>s that are <column name>s of both $T1$ and $T2$ in the order that those <column name>s appear in $T1$.
 - If <corresponding column list> is specified, then let SL be a <select list> of those <column name>s explicitly appearing in the <corresponding column list> in the order that these <column name>s appear in the <corresponding column list>. Every <column name> in the <corresponding column list> shall be a <column name> of both $T1$ and $T2$.
 - The <query term> or <query expression body> is equivalent to:

(SELECT SL FROM $TN1$) OP (SELECT SL FROM $TN2$)

- If CORRESPONDING is not specified, then $T1$ and $T2$ shall be of the same degree.
- 17) If a <query term> is a <query primary>, then the declared type of the <query term> is that of the <query primary>. The column descriptor of the i -th column of the <query term> is the same as the column descriptor of the i -th column of the <query primary>.
- 18) If a <query term> immediately contains INTERSECT, then:
- Let C be the <column name> of the i -th column of $T1$. If the <column name> of the i -th column of $T2$ is C , then the <column name> of the i -th column of TR is C ; otherwise, the <column name> of the i -th column of TR is implementation-dependent.
 - The declared type of the i -th column of TR is determined by applying [Subclause 9.3, “Result of data type combinations”](#), to the declared types of the i -th column of $T1$ and the i -th column of $T2$.
 - If the SQL-implementation supports Feature T101, “Enhanced nullability determination” and the i -th columns of either $T1$ or $T2$ are known not nullable, then the i -th column of TR is known not nullable; otherwise, the i -th column of TR is possibly nullable.

- 19) If a <query term> is a <query primary>, then the column descriptors of the <query term> are the same as the column descriptors of the <query primary>.
- 20) Case:
- a) If a <query expression body> is a <query term>, then the column descriptors of the <query expression body> are the same as the column descriptors of the <query term>.
 - b) If a <query expression body> immediately contains UNION or EXCEPT, then:
 - i) Let C be the <column name> of the i -th column of $T1$. If the <column name> of the i -th column of $T2$ is C , then the <column name> of the i -th column of TR is C ; otherwise, the <column name> of the i -th column of TR is implementation-dependent.
 - ii) If TR is not the result of an anchor expression, then the declared type of the i -th column of TR is determined by applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the i -th column of $T1$ and the i -th column of $T2$.
 - iii) Case:
 - 1) If the SQL-implementation does not support Feature T101, “Enhanced nullability determination”, then the i -th column of TR is possibly nullable.
 - 2) If the <query expression body> immediately contains EXCEPT, then if the i -th column of $T1$ is known not nullable, then the i -th column of TR is known not nullable; otherwise, the i -th column of TR is possibly nullable.
 - 3) Otherwise, if the i -th columns of both $T1$ and $T2$ are known not nullable, then the i -th column of TR is known not nullable; otherwise, the i -th column of TR is possibly nullable.
 - iv) If TR is the result of an anchor expression ARE , then:
 - 1) Let l be the number of recursive tables that belong to the stratum of ARE . For j ranging from 1 (one) to l , let WQT_j be those tables. Of the operands $T1$ and $T2$ of TR , let $TNREC$ be the operand that is the result of the non-recursive operand of ARE and let $TREC$ be the other operand. The i -th column of TR is said to be *recursively referred to* if there exists at least one k , $1 \leq k \leq l$, such that a column of WQT_k is an underlying column of the i -th column of $TREC$. Otherwise, that column is said to be *not recursively referred to*.
 - 2) If the i -th column of TR is not recursively referred to, then the declared type of the i -th column of TR is determined by applying Subclause 9.3, “Result of data type combinations”, to the declared types of the i -th column of $T1$ and the i -th column of $T2$. If the i -th columns of either $T1$ or $T2$ are known not nullable, then the i -th column of TR is *known not nullable*; otherwise, the i -th column of TR is *possibly nullable*.
 - 3) If the i -th column of TR is recursively referred to, then:
 - A) The i -th column of TR is *possibly nullable*.
 - B) Case:
 - I) If $T1$ is $TNREC$, then if the i -th column of TR is recursively referred to, then the declared type of the i -th column of TR is the same as the declared type of the i -th column of $T1$.

- II) If $T2$ is *TNREC*, then if the i -th column of TR is recursively referred to, then the declared type of the i -th column of TR is the same as the declared type of the i -th column of $T2$.
- 21) The *simply underlying tables* of QE are the <table or query name>s, <query specification>s, and <derived table>s contained, without an intervening <derived table> or an intervening <join condition>, in the <query expression body> immediately contained in QE .
- 22) An <explicit table> is *possibly non-deterministic* if the simply contained <table or query name> identifies a viewed table whose original <query expression> is possibly non-deterministic.
- 23) A <query expression> is *possibly non-deterministic* if any of the following are true:
- a) The <query expression> contains a <fetch first clause>.
 - b) The <query expression> is a <query primary> that is possibly non-deterministic.
 - c) UNION, EXCEPT, or INTERSECT is specified and either of the first or second operands is possibly non-deterministic.
 - d) UNION, EXCEPT, or INTERSECT is specified and there is a column of the result such that the declared types $DT1$ and $DT2$ of the column in the two operands have corresponding constituents such that one constituent is datetime with time zone and the other is datetime without time zone.
 - e) Both of the following are true:
 - i) T contains a set operator UNION and ALL is not specified, or T contains either of the set operators EXCEPT or INTERSECT.
 - ii) Exactly one of the following is true:
 - 1) The first or second operand contains a column that has a declared type of character string.
 - 2) The first or second operand contains a column that has a declared type of datetime with time zone.
 - 3) The first or second operand contains a column that has a declared type that is a user-defined type.
- 24) The *underlying columns* of each column of QE and of QE itself are defined as follows:
- a) A column of a <table value constructor> has no underlying columns.
 - b) The underlying columns of every i -th column of a <simple table> ST are the underlying columns of the i -th column of the table immediately contained in ST .
 - c) If no set operator is specified, then the underlying columns of every i -th column of QE are the underlying columns of the i -th column of the <simple table> simply contained in QE .
 - d) If a set operator is specified, then the underlying columns of every i -th column of QE are the underlying columns of the i -th column of $T1$ and those of the i -th column of $T2$.
 - e) Let C be some column. C is an underlying column of QE if and only if C is an underlying column of some column of QE .
- 25) The *updatable columns* of QE are defined as follows:
- a) A column of a <table value constructor> is not an updatable column.

- b) A column of a <simple table> *ST* is an *updatable column* of *ST* if the underlying column of *ST* is updatable.
- c) If no set operator is specified, then a column of *QE* is an *updatable column* of *QE* if its underlying column is updatable.
- d) If a set operator is specified, then

Case:

- i) If the SQL implementation supports Feature T111, “Updatable joins, unions, and columns”, a set operator UNION ALL is specified, and both underlying columns of the *i*-th column of *QE* are updatable, then the *i*-th column of *QE* is an updatable column of *QE*.
- ii) Otherwise, the *i*-th column of *QE* is not updatable.

NOTE 199 — If a set operator UNION DISTINCT, EXCEPT, or INTERSECT is specified, or if the SQL implementation does not support Feature T111, “Updatable joins, unions and columns”, then there are no updatable columns.

26) A <query expression> *QE* shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.

27) If an <order by clause> is specified, then:

- a) Let *OBC* be the <order by clause>. Let *NSK* be the number of <sort specification>s in *OBC*. For each *i* between 1 (one) and *NSK*, let *K_i* be the <sort key> contained in the *i*-th <sort specification> in *OBC*.
- b) Each *K_i* shall contain a <column reference> and shall not contain a <query expression> or a <set function specification>.
- c) If *QE* is a <query expression body> that is a <query term> that is a <query primary> that is a <simple table> that is a <query specification>, then the <query expression> is said to be a *simple table query*.

d) Case:

- i) If <sort specification list> contains any <sort key> *K_i* that contains a column reference to a column that is not a column of *T*, then:
 - 1) *QE* shall be a simple table query.
 - 2) Let *TE* be the <table expression> immediately contained in the <query specification> *QS* contained in *QE*.
 - 3) Let *SL* be the <select list> of *QS*. Let *SLT* be obtained from *SL* by replacing each <column reference> with its fully qualified equivalent.
 - 4) Let *OBCT* be obtained from *OBC* by replacing each <column reference> that references a column of *TE* with its fully qualified equivalent; in the case of common column names, each common column name is regarded as fully qualified.
 - 5) For each *i* between 1 (one) and *NSK*, let *KT_i* be the <sort key> contained in the *i*-th <sort specification> contained in *OBCT*.
 - 6) For each *i* between 1 (one) and *NSK*, if *KT_i* has the same left normal form derivation as the <value expression> immediately contained in some <derived column> *DC* of *SLT*, then:

NOTE 200 — “Left normal form derivation” is defined in Subclause 6.2, “Notation provided in this International Standard”, in [ISO9075-1].

A) Case:

- I) If *DC* simply contains an <as clause>, then let *CN* be the <column name> contained in the <as clause>.
- II) Otherwise, let *CN* be an implementation-dependent <column name> that is not equivalent to the explicit or implicit <column name> of any other <derived column> contained in *SLT*. Let *VE* be the <value expression> simply contained in *DC*. *DC* is replaced in *SLT* by

VE AS CN

B) *KT_i* is replaced in *OBCT* by

CN

7) Let *SCR* be the set of <column reference>s to columns of *TE* that remain in *OBCT* after the preceding transformation.

8) Let *NSCR* be the number of <column reference>s contained in *SCR*. For each *j* between 1 (one) and *NSCR*, let *C_j* be an enumeration of these <column reference>s.

9) Case:

A) If *NSCR* is 0 (zero), then let *SKL* be the zero-length string.

B) Otherwise:

- I) *T* shall not be a grouped table.
- II) *QS* shall not specify the <set quantifier> DISTINCT or directly contain one or more <set function specification>s.
- III) Let *SKL* be the comma-separated list of <derived column>s:

, C₁, C₂, . . . , C_{NSCR}

The columns *C_j* are said to be *extended sort key columns*.

10) Let *SORT* be the table specified by the <query specification>:

SELECT *SLT SKL TE*

11) Let *EOBC* be *OBCT*.

ii) Otherwise, let *SORT* be *T* and let *EOBC* be *OBC*.

e) *SORT* is said to be the *sort table* of *QE*. *EOBC* is said to be the *extended order by clause* of *SORT*.

28) If <fetch first clause> is specified, then

Case:

- a) If <unsigned integer> is specified, then let *UI* be that <unsigned integer>. *UI* shall be greater than 0 (zero).
- b) Otherwise, let *UI* be 1 (one).

Access Rules

None.

General Rules

- 1) If a non-recursive <with clause> is specified, then:
 - a) For every <with list element> *WLE*, let *WQN* be the <query name> immediately contained in *WLE*. Let *WQE* be the <query expression> simply contained in *WLE*. Let *WLT* be the table resulting from evaluation of *WQE*, with each column name replaced by the corresponding element of the <with column list>, if any, immediately contained in *WLE*.
 - b) Every <table reference> contained in <query expression> that specifies *WQN* identifies *WLT*.
- 2) If a potentially recursive <with clause> *WC* is specified, then:
 - a) Let *n* be the number of <with list element>s *WLE_i* of the <with list> *WL* immediately contained in *WC*. For *i* ranging from 1 (one) to *n*, let *WQN_i* and *WQE_i* be the <query name>s and the <query expression>s simply contained in *WLE_i*. Let *WLP_j* be the elements of a partitioning of *WL* such that each *WLP_j* contains all *WLE_i* that belong to one stratum, and let *m* be the number of partitions. Let the *partition dependency graph PDG* of *WL* be a directed graph such that:
 - i) Each partition *WLP_j* of *WL* is represented by exactly one node of *PDG*.
 - ii) There is an arc from the node representing *WLP_j* to the node representing *WLP_k* if and only if *WLP_j* contains at least one *WLE_i*, *WLP_k* contains at least one *WLE_h*, and *WQE_i* contains a <query name> referencing *WQN_h*.
 - b) While the set of nodes of *PDG* is not empty, do:
 - i) Evaluate the partitions of *PDG* that have no outgoing arc.
 - ii) Remove the partitions and their incoming arcs from *PDG*.
 - c) Let *LIP* be some partition of *WL*. Let *m* be the number of <with list element>s in *LIP*, and for *i* ranging from 1 (one) to *m*, let *WLE_i* be a <with list element> of *LIP*, and let *WQN_i* and *WQE_i* be the <query name> and <query expression> simply contained in *WLE_i*. Let *SQE_i* be the set of <query expression>s contained in *WQE_i*. Let *SQE* be a set of <query expression>s such that a <query expression> belongs to *SQE* if and only if it is contained in some *WQE_i*. Let *p* be the number of <query expression>s in *SQE* and let *AQE_k*, $1 \leq k \leq p$ be the *k*-th <query expression> belonging to *SQE*.
 - i) Every <query expression> *AQE_k* that contains a recursive query name in scope is marked as *recursive*.

- ii) Let RT_k and WT_k be tables whose row type is the row type of AQE_k . Let RT_k and WT_k be initially empty. RT_k and WT_k are said to be *associated with* AQE_k . If AQE_k is immediately contained in some WQE_i , then RT_k and WT_k are said to be the *intermediate result table* and *working table*, respectively, *associated with* the <query name> WQN_i .
- iii) If a <query expression> AQE_k not marked as recursive is immediately contained in a <query expression body> that is marked as recursive and that specifies UNION, then AQE_k is marked as *iteration ignorable*.
- iv) For each AQE_k ,

Case:

- 1) If AQE_k consists of a <query specification> that immediately contains DISTINCT, then AQE_k *suppresses duplicates*.
- 2) If AQE_k consists of a <query expression body> or <query term> that explicitly or implicitly immediately contains DISTINCT, then AQE_k *suppresses duplicates*.
- 3) Otherwise, AQE_k does not suppress duplicates.
- v) If an AQE_k is not marked as recursive, then let RT_k and WT_k be the result of AQE_k .
- vi) For every RT_k , let RTN_k be the name of RT_k . If AQE_k is not marked as recursive, then replace AQE_k with:

TABLE RTN_k

- vii) For every WQE_i of LIP , let the *recursive query names in scope* denote the associated result tables. Evaluate every WQE_i . For every AQE_k contained in any such WQE_i , let RT_k and WT_k be the result of AQE_k .

NOTE 201 — This ends the initialization phase of the evaluation of a partition.

- viii) For every AQE_k of LIP that is marked as iteration ignorable, let RT_k be an empty table.
- ix) While some WT_k of LIP is not empty, do:

- 1) Let the recursive query names in scope of LIP denote the associated working tables.
- 2) Evaluate every WQE_i of LIP .
- 3) For every AQE_k that is marked as recursive,

Case:

- A) If AQE_k suppresses duplicates, then let WT_k be the result of AQE_k EXCEPT RTN_k .
- B) Otherwise, let WT_k be the result of AQE_k .
- 4) For every WT_k , let WTN_k be the table name of WT_k . Let RT_k be the result of:

TABLE WTN_k UNION ALL TABLE RTN_k

x) Any reference to WQN_i identifies the intermediate result table RT_k associated with WQN_i .

3) Case:

a) If no set operator is specified, then T is the result of the specified <simple table>.

b) Otherwise:

i) Let D be the degree of T .

ii) For each i , $1 \text{ (one)} \leq i \leq D$:

1) ¹⁴Let DTC_i be the declared type of the i -th column of T .

2) Let $TCN1_i$ be the effective name for the i -th column of $T1$.

3) Let $TCN2_i$ be the effective name for the i -th column of $T2$.

4) Let $ET1$ be a <query expression> of the form

```
SELECT CAST(TCN11 AS DTC1) ,
       CAST(TCN12 AS DTC2) ,
       ... ,
       CAST(TCN1D AS DTCD)
FROM TN1
```

5) Let $ET2$ be a <query expression> of the form

```
SELECT CAST(TCN21 AS DTC1) ,
       CAST(TCN22 AS DTC2) ,
       ... ,
       CAST(TCN2D AS DTCD)
FROM TN2
```

iii) T contains the following rows:

1) Let R be a row that is a duplicate of some row in $ET1$ or of some row in $ET2$ or both. Let m be the number of duplicates of R in $ET1$ and let n be the number of duplicates of R in $ET2$, where $m \geq 0$ and $n \geq 0$.

2) If DISTINCT is specified or implicit, then

Case:

A) If UNION is specified, then

Case:

I) If $m > 0$ or $n > 0$, then T contains exactly one duplicate of R .

II) Otherwise, T contains no duplicate of R .

B) If EXCEPT is specified, then

Case:

I) If $m > 0$ and $n = 0$, then T contains exactly one duplicate of R .

- II) Otherwise, T contains no duplicate of R .
- C) If INTERSECT is specified, then
 - Case:
 - I) If $m > 0$ and $n > 0$, then T contains exactly one duplicate of R .
 - II) Otherwise, T contains no duplicates of R .
- 3) If ALL is specified, then
 - Case:
 - A) If UNION is specified, then the number of duplicates of R that T contains is $(m + n)$.
 - B) If EXCEPT is specified, then the number of duplicates of R that T contains is the maximum of $(m - n)$ and 0 (zero).
 - C) If INTERSECT is specified, then the number of duplicates of R that T contains is the minimum of m and n .

NOTE 202 — See the General Rules of Subclause 8.2, “<comparison predicate>”.

- 4) Case:
 - a) If EXCEPT is specified and a row R of T is replaced by some row RR , then the row of $T1$ from which R is derived is replaced by RR .
 - b) If INTERSECT is specified, then:
 - i) If a row R is inserted into T , then:
 - 1) If $T1$ does not contain a row whose value equals the value of R , then R is inserted into $T1$.
 - 2) If $T1$ contains a row whose value equals the value of R and no row of T is derived from that row, then R is inserted into $T1$.
 - 3) If $T2$ does not contain a row whose value equals the value of R , then R is inserted into $T2$.
 - 4) If $T2$ contains a row whose value equals the value of R and no row of T is derived from that row, then R is inserted into $T2$.
 - ii) If a row R is replaced by some row RR , then:
 - 1) The row of $T1$ from which R is derived is replaced with RR .
 - 2) The row of $T2$ from which R is derived is replaced with RR .
- 5) The rows of T are ordered as follows:
 - a) If QE does not immediately contain an <order by clause>, then the ordering of rows in T is implementation-dependent.
 - b) If QE immediately contains an <order by clause>, then the ordering of rows in T is the same as the ordering of rows in the sort table of QE and its extended order by clause as determined by the General Rules of Subclause 10.10, “<sort specification list>”.

NOTE 203 — The table specified by QE is effectively the sort table of QE with all extended sort key columns (if any) removed. “extended sort key column” and “extended order by clause” are defined in the Syntax Rules of this Subclause.

- 6) If <fetch first clause> is specified, then:
 - a) Let *OCT* be the cardinality of *T*.
 - b) Let *VUI* be the value of *UI*.
 - c) If *VUI* is less than *OCT*, then rows other than the first *VUI* rows in order as specified by [GR 5](#)) of this Subclause are removed from *T* and *VUI* is the cardinality of *T*.

Conformance Rules

- 1) Without Feature T121, “WITH (excluding RECURSIVE) in query expression”, in conforming SQL language, a <query expression> shall not contain a <with clause>.
- 2) Without Feature T122, “WITH (excluding RECURSIVE) in subquery”, in conforming SQL language, a <query expression> contained in a <query expression> shall not contain a <with clause>.
- 3) Without Feature T131, “Recursive query”, conforming SQL language shall not contain a <query expression> that contains RECURSIVE.
- 4) Without Feature T132, “Recursive query in subquery”, in conforming SQL language, a <query expression> contained in a <query expression> shall not contain RECURSIVE.
- 5) Without Feature F661, “Simple tables”, conforming SQL language shall not contain a <simple table> that immediately contains a <table value constructor> except in an <insert statement>.
- 6) Without Feature F661, “Simple tables”, conforming SQL language shall not contain an <explicit table>.
- 7) Without Feature F302, “INTERSECT table operator”, conforming SQL language shall not contain a <query term> that contains INTERSECT.
- 8) Without Feature F301, “CORRESPONDING in query expressions”, conforming SQL language shall not contain a <query expression> that contains CORRESPONDING.
- 9) Without Feature T551, “Optional key words for default syntax”, conforming SQL language shall not contain UNION DISTINCT, EXCEPT DISTINCT, or INTERSECT DISTINCT.
- 10) Without Feature F304, “EXCEPT ALL table operator”, conforming SQL language shall not contain a <query expression> that contains EXCEPT ALL.

NOTE 204 — If DISTINCT, INTERSECT or EXCEPT is specified, then the Conformance Rules of [Subclause 9.10](#), “Grouping operations”, apply.

- 11) Without Feature F850, “Top-level <order by clause> in <query expression>”, in conforming SQL language, a <query expression> not immediately contained in either an <array value constructor by query> or a <cursor specification> shall not immediately contain an <order by clause>.
- 12) Without Feature F851, “<order by clause> in subqueries”, in conforming SQL language, a <query expression> contained in another <query expression> shall not immediately contain an <order by clause>.
- 13) Without Feature F855, “Nested <order by clause> in <query expression>”, in conforming SQL language, a <query primary> shall not immediately contain an <order by clause>.
- 14) Without Feature F856, “Nested <fetch first clause> in <query expression>”, in conforming SQL language, a <query primary> shall not immediately contain a <fetch first clause>.

- 15) Without Feature F857, “Top-level <fetch first clause> in <query expression>”, in conforming SQL language, a <query expression> shall not immediately contain a <fetch first clause>.
- 16) Without Feature F858, “<fetch first clause> in subqueries”, in conforming SQL language, a <query expression> contained in another <query expression> shall not immediately contain a <fetch first clause>.

7.14 <search or cycle clause>

Function

Specify the generation of ordering and cycle detection information in the result of recursive query expressions.

Format

```
<search or cycle clause> ::=
    <search clause>
  | <cycle clause>
  | <search clause> <cycle clause>

<search clause> ::=
    SEARCH <recursive search order> SET <sequence column>

<recursive search order> ::=
    DEPTH FIRST BY <column name list>
  | BREADTH FIRST BY <column name list>

<sequence column> ::=
    <column name>

<cycle clause> ::=
    CYCLE <cycle column list> SET <cycle mark column> TO <cycle mark value>
    DEFAULT <non-cycle mark value> USING <path column>

<cycle column list> ::=
    <cycle column> [ { <comma> <cycle column> }... ]

<cycle column> ::=
    <column name>

<cycle mark column> ::=
    <column name>

<path column> ::=
    <column name>

<cycle mark value> ::=
    <value expression>

<non-cycle mark value> ::=
    <value expression>
```

Syntax Rules

- 1) Let *WLEC* be an expandable <with list element> immediately containing a <search or cycle clause>.
- 2) Let *WQN* be the <query name>, *WCL* the <with column list>, and *WQE* the <query expression> simply contained in *WLEC*. Let *WQEB* be the <query expression body> immediately contained in *WQE*. Let *OP* be the set operator immediately contained in *WQEB*. Let *TLO* be the <query expression body> that constitutes the first operand of *OP* and let *TRO* be the <query specification> that (necessarily) constitutes the second operand of *OP*.

- a) Let *TROSL* be the <select list> immediately contained in *TRO*. Let *WQNTR* be the <table reference> simply contained in the <from clause> immediately contained in the <table expression> *TROTE* immediately contained in *TRO* such that *WQNTR* immediately contains *WQN*.

Case:

- i) If *WQNTR* simply contains a <correlation name>, then let *WQNCRN* be that <correlation name>.
 - ii) Otherwise, let *WQNCRN* be *WQN*.
- b) Case:
- i) If *WLEC* simply contains a <search clause> *SC*, then let *SQC* be the <sequence column> and *SO* be the <recursive search order> immediately contained in *SC*. Let *CNL* be the <column name list> immediately contained in *SO*.

- 1) *WCL* shall not contain a <column name> that is equivalent to *SQC*.
- 2) Every <column name> of *CNL* shall be equivalent to some <column name> contained in *WCL*. No <column name> shall be contained more than once in *CNL*.

- 3) Case:

A) If *SO* immediately contains DEPTH, then let *SCEX1* be:

WQNCRN.SQC

let *SCEX2* be:

SQC || ARRAY [ROW(*CNL*)]

and let *SCIN* be:

ARRAY [ROW(*CNL*)]

B) If *SO* immediately contains BREADTH, then let *SCEX1* be:

```
( SELECT OC.*
  FROM ( VALUES (WQNCRN.SQC) )
        OC(LEVEL, CNL) )
```

let *SCEX2* be:

ROW(*SQC*.LEVEL + 1, *CNL*)

and let *SCIN* be:

ROW(0, *CNL*)

- ii) If *WLEC* simply contains a <cycle clause> *CC*, then let *CCL* be the <cycle column list>, let *CMC* be the <cycle mark column>, let *CMV* be the <cycle mark value>, let *CMD* be the <non-cycle mark value>, and let *CPA* be the <path column> immediately contained in *CC*.
 - 1) Every <column name> of *CCL* shall be equivalent to some <column name> contained in *WCL*. No <column name> shall be contained more than once in *CCL*.

- 2) *CMC* and *CPA* shall not be equivalent to each other and not equivalent to any <column name> of *WCL*.
- 3) The declared type of *CMV* and *CMD* shall be character string of length 1 (one). *CMV* and *CMD* shall be literals and *CMV* shall not be equal to *CMD*.

- 4) Let *CCEX1* be:

WQNCRN.CMC , *WQNCRN.CPA*

Let *CCEX2* be:

CASE WHEN ROW(*CCL*) IN
(SELECT P.* FROM TABLE(*CPA*) P)

THEN *CMV* ELSE *CMD* END,

CPA || ARRAY [ROW(*CCL*)]

Let *CCIN* be:

CMD , ARRAY [ROW(*CCL*)]

Let *NCCONI* be:

CMC <> *CMV*

iii) Case:

- 1) If *WLEC* simply contains a <search clause> and does not simply contain a <cycle clause>, then let *EWCL* be:

WCL , *SQC*

Let *ETLOSL* be:

WCL , *SCIN*

Let *ETROSL* be:

WCL , *SCEX2*

Let *ETROSL1* be:

TROSL , *SCEX1*

Let *NCCON* be:

TRUE

- 2) If *WLEC* simply contains a <cycle clause> and does not simply contain a <search clause>, then let *EWCL* be:

WCL , *CMC* , *CPA*

Let *ETLOSL* be:

WCL, CCIN

Let *ETROSL* be:

WCL, CCEX2

Let *ETROSL1* be:

TROSL, CCEX1

Let *NCCON* be:

NCCON1

3) If *WLEC* simply contains both a <search clause> and a <cycle clause> *CC*, then:

A) The <column name>s *SQC*, *CMC*, and *CPA* shall not be equivalent to each other.

B) Let *EWCL* be:

WCL, SQC, CMC, CPA

Let *ETLOSL* be:

WCL, SCIN, CCIN

Let *ETROSL* be:

WCL, SCEX2, CCEX2

Let *ETROSL1* be:

TROSL, SCEX1, CCEX1

C) Let *NCCON* be:

NCCON1

c) *WLEC* is equivalent to the expanded <with list element>:

```
WQN(EWCL) AS
( SELECT ETLOSL FROM (TLO) TLOCN(WCL)
  OP
  SELECT ETROSL
  FROM (SELECT ETROSL1 TROTE) TROCRN(EWCL)
  WHERE NCCON
)
```

Access Rules

None.

General Rules

None.

Conformance Rules

None.

7.15 <subquery>

Function

Specify a scalar value, a row, or a table derived from a <query expression>.

Format

```
<scalar subquery> ::=
    <subquery>

<row subquery> ::=
    <subquery>

<table subquery> ::=
    <subquery>

<subquery> ::=
    <left paren> <query expression> <right paren>
```

Syntax Rules

- 1) The degree of a <scalar subquery> shall be 1 (one).
- 2) The degree of a <row subquery> shall be greater than 1 (one).
- 3) Let *QE* be the <query expression> simply contained in <subquery>.
- 4) The declared type of a <scalar subquery> is the declared type of the column of *QE*.
- 5) The declared type of a <row subquery> is a row type consisting of one field for each column of *QE*. The declared type and field name of each field of this row type is the declared type and column name of the corresponding column of *QE*.
- 6) The declared types of the columns of a <table subquery> are the declared types of the respective columns of *QE*.

Access Rules

None.

General Rules

- 1) If no SQL-transaction is active for the SQL-agent, then an SQL-transaction is initiated.
- 2) Let *RS* be a <row subquery>. Let *RRS* be the result of the <query expression> simply contained in *RS*. Let *D* be the degree of *RRS*.

Case:

- a) If the cardinality of *RRS* is greater than 1 (one), then an exception condition is raised: *cardinality violation*.

- b) If the cardinality of RRS is 0 (zero), then the value of the <row subquery> is a row whose degree is D and whose fields are all the null value.
 - c) Otherwise, the value of RS is RRS .
- 3) Let SS be a <scalar subquery>.
- Case:
- a) If the cardinality of SS is greater than 1 (one), then an exception condition is raised: *cardinality violation*.
 - b) If the cardinality of SS is 0 (zero), then the value of the <scalar subquery> is the null value.
 - c) Otherwise, let C be the column of <query expression> simply contained in SS . The value of SS is the value of C in the unique row of the result of the <scalar subquery>.

Conformance Rules

None.

(Blank page)

8 Predicates

8.1 <predicate>

This Subclause is modified by Subclause 8.1, “<predicate>”, in ISO/IEC 9075-14.

Function

Specify a condition that can be evaluated to give a boolean value.

Format

```
14 <predicate> ::=
    <comparison predicate>
  | <between predicate>
  | <in predicate>
  | <like predicate>
  | <similar predicate>
  | <regex like predicate>
  | <null predicate>
  | <quantified comparison predicate>
  | <exists predicate>
  | <unique predicate>
  | <normalized predicate>
  | <match predicate>
  | <overlaps predicate>
  | <distinct predicate>
  | <member predicate>
  | <submultiset predicate>
  | <set predicate>
  | <type predicate>
```

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) 14 The result of a <predicate> is the truth value of the immediately contained <comparison predicate>, <between predicate>, <in predicate>, <like predicate>, <similar predicate>, <regex like predicate>, <null predicate>, <quantified comparison predicate>, <exists predicate>, <unique predicate>, <match predicate>, <overlaps predicate>, <distinct predicate>, <member predicate>, <submultiset predicate>, <set predicate>, or <type predicate>.

Conformance Rules

None.

8.2 <comparison predicate>

This Subclause is modified by [Subclause 7.1](#), “<comparison predicate>”, in ISO/IEC 9075-13.

Function

Specify a comparison of two row values.

Format

```
<comparison predicate> ::=
    <row value predicand> <comparison predicate part 2>

<comparison predicate part 2> ::=
    <comp op> <row value predicand>

<comp op> ::=
    <equals operator>
    | <not equals operator>
    | <less than operator>
    | <greater than operator>
    | <less than or equals operator>
    | <greater than or equals operator>
```

Syntax Rules

- 1) The two <row value predicand>s shall be of the same degree.
- 2) Let *corresponding fields* be fields with the same ordinal position in the two <row value predicand>s.
- 3) The declared types of the corresponding fields of the two <row value predicand>s shall be comparable.
- 4) Let R_x and R_y respectively denote the first and second <row value predicand>s.
- 5) Let N be the number of fields in the declared type of R_x . Let X_i , $1 \text{ (one)} \leq i \leq N$, be the i -th field in the declared type of R_x and let Y_i be the i -th field in the declared type of R_y . For each i :
 - a) Case:
 - i) If <comp op> is <equals operator> or <not equals operator>, then X_i and Y_i are operands of an equality operation. The Syntax Rules of [Subclause 9.9](#), “Equality operations”, apply.
 - ii) Otherwise, X_i and Y_i are operands of an ordering operation. The Syntax Rules of [Subclause 9.12](#), “Ordering operations”, apply.
 - b) Case:
 - i) If the declared types of X_i and Y_i are user-defined types, then let $UDT1$ and $UDT2$ be respectively the declared types of X_i and Y_i . $UDT1$ and $UDT2$ shall be in the same subtype family. $UDT1$ and $UDT2$ shall have comparison types.

NOTE 205 — “Comparison type” is defined in [Subclause 4.7.6](#), “User-defined type comparison and assignment”.

13 NOTE 206 — The comparison form and comparison categories included in the user-defined type descriptors of both *UDT1* and *UDT2* are constrained to be the same and to be the same as those of all their supertypes. If the comparison category is either STATE or RELATIVE, then *UDT1* and *UDT2* are constrained to have the same comparison function; if the comparison category is MAP, they are not constrained to have the same comparison function.

- ii) If the declared types of X_i and Y_i are reference types, then the referenced type of the declared type of X_i and the referenced type of the declared type of Y_i shall have a common supertype.
- iii) If the declared types of X_i and Y_i are collection types in which the declared type of the elements are ET_x and ET_y , respectively, then let $RV1$ and $RV2$ be <value expression>s whose declared types are respectively ET_x and ET_y . The Syntax Rules of this Subclause are applied to:

$$RV1 \text{ <comp op> } RV2$$

- iv) If the declared types of X_i and Y_i are row types, then let $RV1$ and $RV2$ be <value expression>s whose declared types are respectively that of X_i and Y_i . The Syntax Rules of this Subclause are applied to:

$$RV1 \text{ <comp op> } RV2$$

- 6) Let CP be the <comparison predicate> “ $R_x \text{ <comp op> } R_y$ ”. If the declared type of R_x is a user-defined type whose comparison category is RELATIVE, then it is implementation-defined whether the following syntactic transformations are applied; otherwise, the following syntactic transformations are applied.

Case:

- a) If the <comp op> is <not equals operator>, then CP is equivalent to:

$$\text{NOT}(R_x = R_y)$$

- b) If the <comp op> is <greater than operator>, then CP is equivalent to:

$$(R_y < R_x)$$

- c) If the <comp op> is <less than or equals operator>, then CP is equivalent to:

$$\begin{aligned} &(R_x < R_y \\ &\quad \text{OR} \\ &R_y = R_x) \end{aligned}$$

- d) If the <comp op> is <greater than or equals operator>, then CP is equivalent to:

$$\begin{aligned} &(R_y < R_x \\ &\quad \text{OR} \\ &R_y = R_x) \end{aligned}$$

Access Rules

None.

General Rules

- 1) Let XV and YV be two values represented by <value expression>s X and Y , respectively. The result of:

$X \text{ <comp op> } Y$

is determined as follows.

Case:

- a) If either XV or YV is the null value, then

$X \text{ <comp op> } Y$

is Unknown.

- b) Otherwise,

Case:

- i) If the declared types of XV and YV are row types with degree N , then let X_i , $1 \text{ (one)} \leq i \leq N$, denote a <value expression> whose value and declared type is that of the i -th field of XV and let Y_i denote a <value expression> whose value and declared type is that of the i -th field of YV . The result of

$X \text{ <comp op> } Y$

is determined as follows:

- 1) $X = Y$ is True if and only if $X_i = Y_i$ is True for all i .
 - 2) $X < Y$ is True if and only if $X_i = Y_i$ is True for all $i < n$ and $X_n < Y_n$ for some n .
 - 3) $X = Y$ is False if and only if NOT ($X_i = Y_i$) is True for some i .
 - 4) $X < Y$ is False if and only if $X = Y$ is True or $Y < X$ is True.
 - 5) $X \text{ <comp op> } Y$ is Unknown if $X \text{ <comp op> } Y$ is neither True nor False.
- ii) If the declared types of XV and YV are array types and the cardinalities of XV and YV are $N1$ and $N2$, respectively, then let X_i , $1 \text{ (one)} \leq i \leq N1$, denote a <value expression> whose value and declared type is that of the i -th element of XV and let Y_i denote a <value expression> whose value and declared type is that of the i -th element of YV . The result of

$X \text{ <comp op> } Y$

is determined as follows:

- 1) $X = Y$ is True if $N1 = 0$ (zero) and $N2 = 0$ (zero).
- 2) $X = Y$ is True if $N1 = N2$ and, for all i , $X_i = Y_i$ is True.
- 3) $X = Y$ is False if and only if $N1 \neq N2$ or NOT ($X_i = Y_i$) is True, for some i .
- 4) $X \text{ <comp op> } Y$ is Unknown if $X \text{ <comp op> } Y$ is neither True nor False.

8.2 <comparison predicate>

- iii) If the declared types of XV and YV are multiset types and the cardinalities of XV and YV are $N1$ and $N2$, respectively, then the result of

$$X \text{ <comp op> } Y$$

is determined as follows.

Case:

- 1) $X = Y$ is True if $N1 = 0$ (zero) and $N2 = 0$ (zero).
 - 2) $X = Y$ is True if $N1 = N2$, and there exist an enumeration XVE_j , $1 \text{ (one)} \leq j \leq N1$, of the elements of XV and an enumeration YVE_j , $1 \text{ (one)} \leq j \leq N1$, of the elements of YV such that for all j , $XVE_j = YVE_j$.
 - 3) $X = Y$ is Unknown if $N1 = N2$, and there exist an enumeration XVE_j , $1 \text{ (one)} \leq j \leq N1$, of the elements of XV and an enumeration YVE_j , $1 \text{ (one)} \leq j \leq N1$, of the elements of YV such that for all j , " $XVE_j = YVE_j$ " is either True or Unknown.
 - 4) Otherwise, $X = Y$ is False.
- iv) If the declared types of XV and YV are user-defined types, then let UDT_x and UDT_y be respectively the declared types of XV and YV . The result of

$$X \text{ <comp op> } Y$$

is determined as follows:

- 1) If the comparison category of UDT_x is MAP, then let $HF1$ be the <routine name> with explicit <schema name> of the comparison function of UDT_x and let $HF2$ be the <routine name> with explicit <schema name> of the comparison function of UDT_y . If $HF1$ identifies an SQL-invoked method, then let HFX be $X.HF1$; otherwise, let HFX be $HF1(X)$. If $HF2$ identifies an SQL-invoked method, then let HFY be $Y.HF2$; otherwise, let HFY be $HF2(Y)$.

$$X \text{ <comp op> } Y$$

has the same result as

$$HFX \text{ <comp op> } HFY$$

- 2) If the comparison category of UDT_x is RELATIVE, then:
 - A) Let RF be the <routine name> with explicit <schema name> of the comparison function of UDT_x .
 - B) $X = Y$

has the same result as

$$RF (X, Y) = 0$$
 - C) $X < Y$

has the same result as

$$RF(X, Y) = -1$$

D) $X <> Y$

has the same result as

$$RF(X, Y) <> 0$$

E) It is implementation-dependent whether $X > Y$

has the same result as

$$RF(X, Y) = 1$$

or has the same result as

$$RF(Y, X) = -1$$

F) $X \leq Y$

has the same result as

$$RF(X, Y) = -1 \text{ OR } RF(X, Y) = 0$$

G) It is implementation-dependent whether $X \geq Y$

has the same result as

$$RF(X, Y) = 1 \text{ OR } RF(X, Y) = 0$$

or has the same result as

$$RF(Y, X) = -1 \text{ OR } RF(X, Y) = 0$$

NOTE 207 — Since it is implementation-defined whether to use the syntactic transformation to define all comparisons in terms of = and <, for portability, the application should ensure that $RF(X, Y) = -RF(Y, X)$ for all X and Y.

3) 13 If the comparison category of UDT_x is STATE, then:

A) Let SF be the <routine name> of the comparison function of UDT_x .

B) $X = Y$

has the same result as

$$SF(X, Y) = \text{TRUE}$$

v) Otherwise, the result of

$$X \text{ <comp op> } Y$$

is True or False as follows:

1) $X = Y$

is True if and only if XV and YV are equal.

2) $X < Y$

is True if and only if XV is less than YV .

3) $X <_{\text{comp op}} Y$

is False if and only if

$X <_{\text{comp op}} Y$

is not True

2) Numbers are compared with respect to their algebraic value.

3) The comparison of two character strings is determined as follows:

- a) Syntax Rules of [Subclause 9.13, “Collation determination”](#), are applied with the declared types of the two character strings as [VALSET](#), yielding *CS* returned from [COLL](#).
- b) If the length in characters of X is not equal to the length in characters of Y , then the shorter string is effectively replaced, for the purposes of comparison, with a copy of itself that has been extended to the length of the longer string by concatenation on the right of one or more pad characters, where the pad character is chosen based on *CS*. If *CS* has the NO PAD characteristic, then the pad character is an implementation-dependent character different from any character in the character set of X and Y that collates less than any string under *CS*. Otherwise, the pad character is a <space>.
- c) The result of the comparison of X and Y is given by the collation *CS*.
- d) Depending on the collation, two strings may compare as equal even if they are of different lengths or contain different sequences of characters. When any of the operations MAX, MIN, and DISTINCT reference a grouping column, and the UNION, EXCEPT, and INTERSECT operators refer to character strings, the specific value selected by these operations from a set of such equal values is implementation-dependent.

4) The comparison of two binary string values, X and Y , of which at least one is a binary large object string value, is determined by comparison of their octets with the same ordinal position. Let X_i and Y_i be the values of the i -th octets of X and Y , respectively, and let L_x be the length in octets of X and let L_y be the length in octets of Y . X is equal to Y if and only if $L_x = L_y$ and if $X_i = Y_i$ for all i .

5) The comparison of two binary string values X and Y , neither of which is a binary large object string value, is determined as follows:

- a) Let L_x be the length in octets of X and let L_y be the length in octets of Y . Let X_i , $1 \text{ (one)} \leq i \leq L_x$, be the value of the i -th octet of X , and let Y_i , $1 \text{ (one)} \leq i \leq L_y$, be the value of the i -th octet of Y .
- b) If $L_x = L_y$ and $X_i = Y_i$, $1 \text{ (one)} \leq i \leq L_x$, then X is equal to Y .
- c) If $L_x < L_y$, $X_i = Y_i$ for all $i \leq L_x$, and the right-most $L_y - L_x$ octets of Y are all X'00's, then it is implementation-defined whether X is equal to Y or whether X is less than Y .
- d) If $L_x < L_y$, $X_i = Y_i$ for all $i \leq L_x$, and at least one of the right-most $L_y - L_x$ octets of Y is not X'00', then X is less than Y .
- e) If $X_j < Y_j$, for some j , $0 \text{ (zero)} < j \leq \text{minimum}(L_x, L_y)$, and $X_i = Y_i$ for all $i < j$, then X is less than Y .

- 6) The comparison of two datetimes is determined according to the interval resulting from their subtraction. Let X and Y be the two values to be compared and let H be the least significant <primary datetime field> of X and Y , including fractional seconds precision if the data type is time or timestamp.

- a) X is equal to Y if and only if

$(X - Y) \text{ INTERVAL } H = \text{INTERVAL '0' } H$

is True.

- b) X is less than Y if and only if

$(X - Y) \text{ INTERVAL } H < \text{INTERVAL '0' } H$

is True.

NOTE 208 — Two datetimes are comparable only if they have the same <primary datetime field>s; see Subclause 4.6.2, “Datetimes”.

- 7) The comparison of two intervals is determined by the comparison of their corresponding values after conversion to integers in some common base unit. Let X and Y be the two intervals to be compared. Let $A \text{ TO } B$ be the specified or implied interval qualifier of X and $C \text{ TO } D$ be the specified or implied interval qualifier of Y . Let T be the least significant <primary datetime field> of B and D and let U be an interval qualifier of the form $T(N)$, where N is an <interval leading field precision> large enough so that significance is not lost in the CAST operation.

Let XVE be the <value expression>

$\text{CAST} (X \text{ AS INTERVAL } U)$

Let YVE be the <value expression>

$\text{CAST} (Y \text{ AS INTERVAL } U)$

- a) X is equal to Y if and only if

$\text{CAST} (XVE \text{ AS INTEGER }) = \text{CAST} (YVE \text{ AS INTEGER })$

is True.

- b) X is less than Y if and only if

$\text{CAST} (XVE \text{ AS INTEGER }) < \text{CAST} (YVE \text{ AS INTEGER })$

is True.

- 8) In comparisons of boolean values, True is greater than False
- 9) The result of comparing two reference values X and Y is determined by the comparison of their octets with the same ordinal position. Let L_x be the length in octets of X and let L_y be the length in octets of Y . Let X_i and Y_i , $1 \text{ (one)} \leq i \leq L_x$, be the values of the i -th octets of X and Y , respectively. X is equal to Y if and only if $L_x = L_y$ and, for all i , $X_i = Y_i$.

Conformance Rules

None.

NOTE 209 — If <comp op> is <equals operator> or <not equals operator>, then the Conformance Rules of [Subclause 9.9](#), “Equality operations”, apply. Otherwise, the Conformance Rules of [Subclause 9.12](#), “Ordering operations”, apply.

8.3 <between predicate>

Function

Specify a range comparison.

Format

```
<between predicate> ::=  
    <row value predicand> <between predicate part 2>  
  
<between predicate part 2> ::=  
    [ NOT ] BETWEEN [ ASYMMETRIC | SYMMETRIC ]  
        <row value predicand> AND <row value predicand>
```

Syntax Rules

- 1) If neither SYMMETRIC nor ASYMMETRIC is specified, then ASYMMETRIC is implicit.
- 2) Let X, Y, and Z be the first, second, and third <row value predicand>s, respectively.
- 3) “X NOT BETWEEN SYMMETRIC Y AND Z” is equivalent to “NOT (X BETWEEN SYMMETRIC Y AND Z)”.
- 4) “X BETWEEN SYMMETRIC Y AND Z” is equivalent to “((X BETWEEN ASYMMETRIC Y AND Z) OR (X BETWEEN ASYMMETRIC Z AND Y))”.
- 5) “X NOT BETWEEN ASYMMETRIC Y AND Z” is equivalent to “NOT (X BETWEEN ASYMMETRIC Y AND Z)”.
- 6) “X BETWEEN ASYMMETRIC Y AND Z” is equivalent to “X>=Y AND X<=Z”.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature T461, “Symmetric BETWEEN predicate”, conforming SQL language shall not contain SYMMETRIC or ASYMMETRIC.

NOTE 210 — Since <between predicate> is an ordering operation, the Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.

8.4 <in predicate>

Function

Specify a quantified comparison.

Format

```

<in predicate> ::=
    <row value predicand> <in predicate part 2>

<in predicate part 2> ::=
    [ NOT ] IN <in predicate value>

<in predicate value> ::=
    <table subquery>
    | <left paren> <in value list> <right paren>

<in value list> ::=
    <row value expression> [ { <comma> <row value expression> }... ]

```

Syntax Rules

- 1) If <in value list> consists of a single <row value expression>, then that <row value expression> shall not be a <scalar subquery>.

NOTE 211 — This Syntax Rule resolves an ambiguity in which <in predicate value> might be interpreted either as a <table subquery> or as a <scalar subquery>. The ambiguity is resolved by adopting the interpretation that the <in predicate value> will be interpreted as a <table subquery>.

- 2) Let *IVL* be an <in value list>.

(*IVL*)

is equivalent to the <table value constructor>:

(VALUES *IVL*)

- 3) Let *RVC* be the <row value predicand> and let *IPV* be the <in predicate value>.

- 4) The expression

RVC NOT IN *IPV*

is equivalent to

NOT (*RVC* IN *IPV*)

- 5) The expression

RVC IN *IPV*

is equivalent to

RVC = ANY IPV

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature F561, “Full value expressions”, conforming SQL language shall not contain a <row value expression> immediately contained in an <in value list> that is not a <value specification>.

NOTE 212 — Since <in predicate> is an equality operation, the Conformance Rules of Subclause 9.9, “Equality operations”, also apply.

8.5 <like predicate>

Function

Specify a pattern-match comparison.

Format

```

<like predicate> ::=
    <character like predicate>
  | <octet like predicate>

<character like predicate> ::=
    <row value predicand> <character like predicate part 2>

<character like predicate part 2> ::=
    [ NOT ] LIKE <character pattern> [ ESCAPE <escape character> ]

<character pattern> ::=
    <character value expression>

<escape character> ::=
    <character value expression>

<octet like predicate> ::=
    <row value predicand> <octet like predicate part 2>

<octet like predicate part 2> ::=
    [ NOT ] LIKE <octet pattern> [ ESCAPE <escape octet> ]

<octet pattern> ::=
    <binary value expression>

<escape octet> ::=
    <binary value expression>

```

Syntax Rules

- 1) The <row value predicand> immediately contained in <character like predicate> shall be a <row value constructor predicand> that is a <common value expression> *CVE*. The declared types of *CVE*, <character pattern>, and <escape character> shall be character string. *CVE*, <character pattern>, and <escape character> shall be comparable.
- 2) The <row value predicand> immediately contained in <octet like predicate> shall be a <row value constructor element> that is a <common value expression> *OVE*. The declared types of *OVE*, <octet pattern>, and <escape octet> shall be binary string.
- 3) If <character like predicate> is specified, then:
 - a) Let *MC* be the <character value expression> of *CVE*, let *PC* be the <character value expression> of the <character pattern>, and let *EC* be the <character value expression> of the <escape character> if one is specified.
 - b) *MC* NOT LIKE *PC*

is equivalent to

NOT (MC LIKE PC)

- c) MC NOT LIKE PC ESCAPE EC

is equivalent to

NOT (MC LIKE PC ESCAPE EC)

- d) The Syntax Rules of [Subclause 9.13, “Collation determination”](#), are applied with *CVE*, *PC*, and (if specified) *EC* as *VALSET*, yielding the collation used for <like predicate> returned from *COLL*.

It is implementation-defined which collations can be used as collations for the <like predicate>.

- 4) If <octet like predicate> is specified, then:

- a) Let *MB* be the <binary value expression> of the *OVE*, let *PB* be the <binary value expression> of the <octet pattern>, and let *EB* be the <binary value expression> of the <escape octet> if one is specified.

- b) MB NOT LIKE PB

is equivalent to

NOT (MB LIKE PB)

- c) MB NOT LIKE PB ESCAPE EB

is equivalent to

NOT (MB LIKE PB ESCAPE EB)

Access Rules

None.

General Rules

- 1) Let *MCV* be the value of *MC* and let *PCV* be the value of *PC*. If *EC* is specified, then let *ECV* be its value.
- 2) Let *MBV* be the value of *MB* and let *PBV* be the value of *PB*. If *EB* is specified, then let *EBV* be its value.
- 3) If <character like predicate> is specified, then:

- a) Case:

- i) If ESCAPE is not specified and either *MCV* or *PCV* are null values, then the result of

MC LIKE PC

is *Unknown*.

- ii) If ESCAPE is specified and one or more of *MCV*, *PCV* and *ECV* are null values, then the result of

MC LIKE PC ESCAPE EC

is *Unknown*.

NOTE 213 — If none of *MCV*, *PCV*, and *ECV* (if present) are null values, then the result is either *True* or *False*.

b) Case:

i) If an <escape character> is specified, then:

- 1) If the length in characters of *ECV* is not equal to 1, then an exception condition is raised: *data exception — invalid escape character*.
- 2) If there is not a partitioning of the string *PCV* into substrings such that each substring has length 1 (one) or 2, no substring of length 1 (one) is the escape character *ECV*, and each substring of length 2 is the escape character *ECV* followed by either the escape character *ECV*, an <underscore> character, or the <percent> character, then an exception condition is raised: *data exception — invalid escape sequence*.

If there is such a partitioning of *PCV*, then in that partitioning, each substring with length 2 represents a single occurrence of the second character of that substring and is called a *single character specifier*.

Each substring with length 1 (one) that is the <underscore> character represents an *arbitrary character specifier*. Each substring with length 1 (one) that is the <percent> character represents an *arbitrary string specifier*. Each substring with length 1 (one) that is neither the <underscore> character nor the <percent> character represents the character that it contains and is called a *single character specifier*.

ii) If an <escape character> is not specified, then each <underscore> character in *PCV* represents an arbitrary character specifier, each <percent> character in *PCV* represents an arbitrary string specifier, and each character in *PCV* that is neither the <underscore> character nor the <percent> character represents itself and is called a *single character specifier*.

c) The string *PCV* is a sequence of the minimum number of substring specifiers such that each portion of *PCV* is part of exactly one substring specifier. A *substring specifier* is an arbitrary character specifier, arbitrary string specifier, or any sequence of single character specifiers.

d) Case:

i) If *MCV* and *PCV* are character strings whose lengths are variable and if the lengths of both *MCV* and *PCV* are 0 (zero), then

MC LIKE PC

is *True*.

ii) The <predicate>

MC LIKE PC

is *True* if there exists a partitioning of *MCV* into substrings such that:

- 1) A substring of *MCV* is a sequence of 0 (zero) or more contiguous characters of *MCV* and each character of *MCV* is part of exactly one substring.

- 2) If the *i*-th substring of *PCV* is an arbitrary character specifier, then the *i*-th substring of *MCV* is any single character.
 - 3) If the *i*-th substring of *PCV* is an arbitrary string specifier, then the *i*-th substring of *MCV* is any sequence of 0 (zero) or more characters.
 - 4) If the *i*-th substring of *PCV* is a single character specifier, then the *i*-th substring of *MCV* contains exactly 1 (one) character that is equal to the character represented by the single character specifier according to the collation of the <like predicate>.
 - 5) The number of substrings of *MCV* is equal to the number of substring specifiers of *PCV*.
- iii) Otherwise,
- MC* LIKE *PC*
- is False.

4) If <octet like predicate> is specified, then:

- a) Case:
 - i) If ESCAPE is not specified and either *MBV* or *PBV* are null values, then the result of

MB LIKE *PB*

 is Unknown.
 - ii) If ESCAPE is specified and one or more of *MBV*, *PBV* and *EBV* are null values, then the result of

MB LIKE *PB* ESCAPE *EB*

 is Unknown.

NOTE 214 — If none of *MBV*, *PBV*, and *EBV* (if present) are null values, then the result is either True or False.

- b) <percent> in the context of an <octet like predicate> has the same bit pattern as the encoding of a <percent> in the SQL_TEXT character set.
- c) <underscore> in the context of an <octet like predicate> has the same bit pattern as the encoding of an <underscore> in the SQL_TEXT character set.
- d) Case:
 - i) If an <escape octet> is specified, then:
 - 1) If the length in octets of *EBV* is not equal to 1, then an exception condition is raised: *data exception — invalid escape octet*.
 - 2) If there is not a partitioning of the string *PBV* into substrings such that each substring has length 1 (one) or 2, no substring of length 1 (one) is the escape octet *EBV*, and each substring of length 2 is the escape octet *EBV* followed by either the escape octet *EBV*, an <underscore> octet, or the <percent> octet, then an exception condition is raised: *data exception — invalid escape sequence*.

If there is such a partitioning of *PBV*, then in that partitioning, each substring with length 2 represents a single occurrence of the second octet of that substring. Each substring with

length 1 (one) that is the <underscore> octet represents an *arbitrary octet specifier*. Each substring with length 1 (one) that is the <percent> octet represents an *arbitrary string specifier*. Each substring with length 1 (one) that is neither the <underscore> octet nor the <percent> octet represents the octet that it contains.

- ii) If an <escape octet> is not specified, then each <underscore> octet in *PBV* represents an arbitrary octet specifier, each <percent> octet in *PBV* represents an arbitrary string specifier, and each octet in *PBV* that is neither the <underscore> octet nor the <percent> octet represents itself.
- e) The string *PBV* is a sequence of the minimum number of substring specifiers such that each portion of *PBV* is part of exactly one substring specifier. A *substring specifier* is an arbitrary octet specifier, an arbitrary string specifier, or any sequence of octets other than an arbitrary octet specifier or an arbitrary string specifier.
- f) Case:

- i) If the lengths of both *MBV* and *PBV* are 0 (zero), then

MB LIKE *PB*

is True.

- ii) The <predicate>

MB LIKE *PB*

is True if there exists a partitioning of *MBV* into substrings such that:

- 1) A substring of *MBV* is a sequence of 0 (zero) or more contiguous octets of *MBV* and each octet of *MBV* is part of exactly one substring.
- 2) If the *i*-th substring specifier of *PBV* is an arbitrary octet specifier, the *i*-th substring of *MBV* is any single octet.
- 3) the *i*-th substring specifier of *PBV* is an arbitrary string specifier, then the *i*-th substring of *MBV* is any sequence of 0 (zero) or more octets.
- 4) If the *i*-th substring specifier of *PBV* is neither an arbitrary character specifier nor an arbitrary string specifier, then the *i*-th substring of *MBV* has the same length and bit pattern as that of the substring specifier.
- 5) The number of substrings of *MBV* is equal to the number of substring specifiers of *PBV*.

- iii) Otherwise:

MB LIKE *PB*

is False.

Conformance Rules

- 1) Without Feature T042, “Extended LOB data type support”, or Feature T022, “Advanced support for BINARY and VARBINARY data types”, conforming SQL language shall not contain an <octet like predicate>.

- 2) Without Feature F281, “LIKE enhancements”, conforming SQL language shall not contain a <common value expression> simply contained in the <row value predicand> immediately contained in <character like predicate> that is not a column reference.
- 3) Without Feature F281, “LIKE enhancements”, conforming SQL language shall not contain a <character pattern> that is not a <value specification>.
- 4) Without Feature F281, “LIKE enhancements”, conforming SQL language shall not contain an <escape character> that is not a <value specification>.
- 5) Without Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <like predicate> shall not be of declared type CHARACTER LARGE OBJECT
- 6) Without Feature F421, “National character”, and Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <like predicate> shall not be of declared type NATIONAL CHARACTER LARGE OBJECT.

8.6 <similar predicate>

Function

Specify a character string similarity by means of a regular expression.

Format

```

<similar predicate> ::=
    <row value predicand> <similar predicate part 2>

<similar predicate part 2> ::=
    [ NOT ] SIMILAR TO <similar pattern> [ ESCAPE <escape character> ]

<similar pattern> ::=
    <character value expression>

<regular expression> ::=
    <regular term>
    | <regular expression> <vertical bar> <regular term>

<regular term> ::=
    <regular factor>
    | <regular term> <regular factor>

<regular factor> ::=
    <regular primary>
    | <regular primary> <asterisk>
    | <regular primary> <plus sign>
    | <regular primary> <question mark>
    | <regular primary> <repeat factor>

<repeat factor> ::=
    <left brace> <low value> [ <upper limit> ] <right brace>

<upper limit> ::=
    <comma> [ <high value> ]

<low value> ::=
    <unsigned integer>

<high value> ::=
    <unsigned integer>

<regular primary> ::=
    <character specifier>
    | <percent>
    | <regular character set>
    | <left paren> <regular expression> <right paren>

<character specifier> ::=
    <non-escaped character>
    | <escaped character>

<non-escaped character> ::=
    !! See the Syntax Rules.

```

```

<escaped character> ::=
    !! See the Syntax Rules.

<regular character set> ::=
    <underscore>
    | <left bracket> <character enumeration>... <right bracket>
    | <left bracket> <circumflex> <character enumeration>... <right bracket>
    | <left bracket> <character enumeration include>...
      <circumflex> <character enumeration exclude>... <right bracket>

<character enumeration include> ::=
    <character enumeration>

<character enumeration exclude> ::=
    <character enumeration>

<character enumeration> ::=
    <character specifier>
    | <character specifier> <minus sign> <character specifier>
    | <left bracket> <colon> <regular character set identifier> <colon> <right bracket>

<regular character set identifier> ::=
    <identifier>

```

Syntax Rules

- 1) The <row value predicand> shall be a <row value constructor predicand> that is a <common value expression> *CVE*. The declared types of *CVE*, <similar pattern>, and <escape character> shall be character string. *CVE*, <similar pattern>, and <escape character> shall be comparable.

NOTE 215 — An implication of the preceding Rule is that *CVE* immediately contains a <string value expression> that, in turn, immediately contains a <character value expression>. In other words, *CVE* is necessarily a <character value expression>.

- 2) Let *CM* be the <character value expression> of *CVE* and let *SP* be the <similar pattern>. If <escape character> *EC* is specified, then

CM NOT SIMILAR TO *SP* ESCAPE *EC*

is equivalent to

NOT (*CM* SIMILAR TO *SP* ESCAPE *EC*)

If <escape character> *EC* is not specified, then

CM NOT SIMILAR TO *SP*

is equivalent to

NOT (*CM* SIMILAR TO *SP*)

- 3) The value of the <identifier> that is a <regular character set identifier> shall be either ALPHA, UPPER, LOWER, DIGIT, ALNUM, SPACE, or WHITESPACE.
- 4) The Syntax Rules of Subclause 9.13, “Collation determination”, are applied with *CVE*, *SP*, and (if specified) *EC* as *VALSET*, yielding the collation used for <similar predicate> returned from *COLL*.

It is implementation-defined which collations can be used as collations for <similar predicate>.

8.6 <similar predicate>

- 5) A <non-escaped character> is any single character from the character set of the <similar pattern> that is not a <left bracket>, <right bracket>, <left paren>, <right paren>, <vertical bar>, <circumflex>, <minus sign>, <plus sign>, <asterisk>, <underscore>, <percent>, <question mark>, <left brace>, or the character specified by the result of the <character value expression> of <escape character>. A <character specifier> that is a <non-escaped character> represents itself.
- 6) An <escaped character> is a sequence of two characters: the character specified by the result of the <character value expression> of <escape character>, followed by a second character that is a <left bracket>, <right bracket>, <left paren>, <right paren>, <vertical bar>, <circumflex>, <minus sign>, <plus sign>, <asterisk>, <underscore>, <percent>, <question mark>, <left brace>, or the character specified by the result of the <character value expression> of <escape character>. A <character specifier> that is an <escaped character> represents its second character.
- 7) The value of <low value> shall be a positive integer. The value of <high value> shall be greater than or equal to the value of <low value>.

Access Rules

None.

General Rules

- 1) Let *MCV* be the result of the <character value expression> of *CVE* and let *PCV* be the result of the <character value expression> of the <similar pattern>. If *EC* is specified, then let *ECV* be its value.
- 2) If the result of the <character value expression> of the <similar pattern> is not a zero-length string and does not have the format of a <regular expression>, then an exception condition is raised: *data exception — invalid regular expression*.
- 3) If an <escape character> is specified, then:
 - a) If the length in characters of *ECV* is not equal to 1 (one), then an exception condition is raised: *data exception — invalid escape character*.
 - b) If *ECV* is one of <left bracket>, <right bracket>, <left paren>, <right paren>, <vertical bar>, <circumflex>, <minus sign>, <plus sign>, <asterisk>, <underscore>, <percent>, <question mark>, or <left brace> and *ECV* occurs in the <regular expression> except in an <escaped character>, then an exception condition is raised: *data exception — invalid use of escape character*.
 - c) If *ECV* is a <colon> and the <regular expression> contains a <regular character set identifier>, then an exception condition is raised: *data exception — escape character conflict*.
- 4) Case:
 - a) If ESCAPE is not specified, then if either or both of *MCV* and *PCV* are the null value, then the result of

CM SIMILAR TO *SP*

 is *Unknown*.
 - b) If ESCAPE is specified, then if one or more of *MCV*, *PCV*, and *ECV* are the null value, then the result of

CM SIMILAR TO SP ESCAPE EC

is *Unknown*.

NOTE 216 — If none of *MCV*, *PCV*, and *ECV* (if present) are the null value, then the result is either *True* or *False*.

- 5) The set of characters in a <character enumeration> is defined as
 - a) If the enumeration is specified in the form “<character specifier> <minus sign> <character specifier>”, then the set of all characters that collate greater than or equal to the character represented by the left <character specifier> and less than or equal to the character represented by the right <character specifier>, according to the collation of the pattern *PCV*.
 - b) Otherwise, the character that the <character specifier> in the <character enumeration> represents.
- 6) Let *LV* be the value of the <low value> contained in a <repeat factor> *RF*.
Case:
 - a) If *RF* does not contain an <upper limit>, then let *HV* be *LV*.
 - b) If *RF* contains an <upper limit> that contains a <high value>, then let *HV* be the value of <high value>.
 - c) Otherwise, let *HV* be the length or maximum length of *CVE*.
- 7) Let *R* be the result of the <character value expression> of the <similar pattern>. The regular language $L(R)$ of the <similar pattern> is a (possibly infinite) set of strings. It is defined recursively for well-formed <regular expression>s *Q*, *Q1*, and *Q2* by the following rules:
 - a) $L(Q1 \text{ <vertical bar> } Q2)$
is the union of $L(Q1)$ and $L(Q2)$
 - b) $L(Q \text{ <asterisk> })$
is the set of all strings that can be constructed by concatenating zero or more strings from $L(Q)$.
 - c) $L(Q \text{ <plus sign> })$
is the set of all strings that can be constructed by concatenating one or more strings from $L(Q)$.
 - d) $L(Q \text{ <repeat factor> })$
is the set of all strings that can be constructed by concatenating *NS*, $LV \leq NS \leq HV$, strings from $L(Q)$.
 - e) $L(\text{<character specifier>})$
is a set that contains a single string of length 1 (one) with the character that the <character specifier> represents
 - f) $L(\text{<percent>})$
is the set of all strings of any length (zero or more) from the character set of the pattern *PCV*.
 - g) $L(Q \text{ <question mark> })$
is the set of all strings that can be constructed by concatenating exactly 0 (zero) or 1 (one) strings from $L(Q)$.

- h) $L(\langle \text{left paren} \rangle Q \langle \text{right paren} \rangle)$
is equal to $L(Q)$
- i) $L(\langle \text{underscore} \rangle)$
is the set of all strings of length 1 (one) from the character set of the pattern *PCV*.
- j) $L(\langle \text{left bracket} \rangle \langle \text{character enumeration} \rangle \langle \text{right bracket} \rangle)$
is the set of all strings of length 1 (one) from the set of characters in the $\langle \text{character enumeration} \rangle$ s.
- k) $L(\langle \text{left bracket} \rangle \langle \text{circumflex} \rangle \langle \text{character enumeration} \rangle \langle \text{right bracket} \rangle)$
is the set of all strings of length 1 (one) with characters from the character set of the pattern *PCV* that are not contained in the set of characters in the $\langle \text{character enumeration} \rangle$.
- l) $L(\langle \text{left bracket} \rangle \langle \text{character enumeration include} \rangle \langle \text{circumflex} \rangle \langle \text{character enumeration exclude} \rangle \langle \text{right bracket} \rangle)$
is the set of all strings of length 1 (one) taken from the set of characters in the $\langle \text{character enumeration include} \rangle$ s, except for those strings of length 1 (one) taken from the set of characters in the $\langle \text{character enumeration exclude} \rangle$.
- m) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{ALPHA} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are $\langle \text{simple Latin letter} \rangle$ s.
- n) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{UPPER} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are $\langle \text{simple Latin upper case letter} \rangle$ s.
- o) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{LOWER} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are $\langle \text{simple Latin lower case letter} \rangle$ s.
- p) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{DIGIT} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are $\langle \text{digit} \rangle$ s.
- q) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{SPACE} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are the $\langle \text{space} \rangle$ character.
- r) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{WHITESPACE} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are white space characters.

NOTE 217 — “white space” is defined in Subclause 3.1.6, “Definitions provided in Part 2”.
- s) $L(\langle \text{left bracket} \rangle \langle \text{colon} \rangle \text{ALNUM} \langle \text{colon} \rangle \langle \text{right bracket} \rangle)$
is the set of all character strings of length 1 (one) that are $\langle \text{simple Latin letter} \rangle$ s or $\langle \text{digit} \rangle$ s.
- t) $L(Q1 \mid Q2)$
is the set of all strings that can be constructed by concatenating one element of $L(Q1)$ and one element of $L(Q2)$.
- u) $L(Q)$

is the set of the zero-length string, if Q is an empty regular expression.

- 8) The <similar predicate>

CM SIMILAR TO *SP*

is True, if there exists at least one element X of $L(R)$ that is equal to MCV according to the collation of the <similar predicate>; otherwise, it is False.

NOTE 218 — The <similar predicate> is defined differently from equivalent forms of the LIKE predicate. In particular, blanks at the end of a pattern and collation are handled differently.

Conformance Rules

- 1) Without Feature T141, “SIMILAR predicate”, conforming SQL language shall not contain a <similar predicate>.
- 2) Without Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <similar predicate> shall not be of declared type CHARACTER LARGE OBJECT.

8.7 <regex like predicate>

Function

Specify a pattern-match comparison using an XQuery regular expression.

Format

```
<regex like predicate> ::=
  <row value predicand> <regex like predicate part 2>

<regex like predicate part 2> ::=
  [ NOT ] LIKE_REGEX <XQuery pattern> [ FLAG <XQuery option flag> ]
```

Syntax Rules

- 1) The <row value predicand> immediately contained in <regex like predicate> shall be a <row value constructor predicand> that is a <common value expression> *CVE* whose declared type is character string.
- 2) If <XQuery option flag> is not specified, then the zero-length character string is implicit.
- 3) Let *PAT* be the <XQuery pattern> and let *FL* be the implicit or explicit <XQuery option flag>.

CVE NOT LIKE_REGEX *PAT* FLAG *FL*

is equivalent to

NOT (*CVE* LIKE_REGEX *PAT* FLAG *FL*)

Access Rules

None.

General Rules

- 1) The value of the <predicate>

CVE LIKE_REGEX *PAT* FLAG *FL*

is

Case:

- a) If the value of *CVE*, *PAT*, or *FL* is the null value, then Unknown.
- b) If the list of match vectors that is the result of applying the General Rules of Subclause 9.16, “XQuery regular expression matching”, with *CVE* as *STRING*, *PAT* as *PATTERN*, 1 (one) as *POSITION*, CHARACTERS as *UNITS*, and *FL* as *FLAG* is non-empty, then True.
- c) Otherwise, False.

Conformance Rules

- 1) Without Feature F841, “LIKE_REGEX predicate”, conforming SQL language shall not contain <regex like predicate>.
- 2) Without Feature F847, “Nonconstant regular expressions”, in conforming SQL language, <XQuery pattern> and <XQuery option flag> shall be <value specification>s.

8.8 <null predicate>

Function

Specify a test for a null value.

Format

```
<null predicate> ::=
  <row value predicand> <null predicate part 2>

<null predicate part 2> ::=
  IS [ NOT ] NULL
```

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let R be the <row value predicand> and let V be the value of R .
- 2) Case:
 - a) If V is the null value, then “ R IS NULL” is True and the value of “ R IS NOT NULL” is False.
 - b) Otherwise:
 - i) The value of “ R IS NULL” is

Case:

 - 1) If the value of every field of V is the null value, then True.
 - 2) Otherwise, False.
 - ii) The value of “ R IS NOT NULL” is

Case:

 - 1) If the value of no field of V is the null value, then True.
 - 2) Otherwise, False.

NOTE 219 — For all R , “ R IS NOT NULL” has the same result as “NOT R IS NULL” if and only if R is of degree 1. Table 14, “<null predicate> semantics”, specifies this behavior.

Table 14 — <null predicate> semantics

Expression	<i>R</i> IS NULL	<i>R</i> IS NOT NULL	NOT <i>R</i> IS NULL	NOT <i>R</i> IS NOT NULL
degree 1: null	<u>True</u>	<u>False</u>	<u>False</u>	<u>True</u>
degree 1: not null	<u>False</u>	<u>True</u>	<u>True</u>	<u>False</u>
degree > 1: all null	<u>True</u>	<u>False</u>	<u>False</u>	<u>True</u>
degree > 1: some null	<u>False</u>	<u>False</u>	<u>True</u>	<u>True</u>
degree > 1: none null	<u>False</u>	<u>True</u>	<u>True</u>	<u>False</u>

Conformance Rules

None.

8.9 <quantified comparison predicate>

Function

Specify a quantified comparison.

Format

```

<quantified comparison predicate> ::=
    <row value predicand> <quantified comparison predicate part 2>

<quantified comparison predicate part 2> ::=
    <comp op> <quantifier> <table subquery>

<quantifier> ::=
    <all>
    | <some>

<all> ::=
    ALL

<some> ::=
    SOME
    | ANY

```

Syntax Rules

- 1) Let *RV1* and *RV2* be <row value predicand>s whose declared types are respectively that of the <row value predicand> and the row type of the <table subquery>. The Syntax Rules of [Subclause 8.2](#), “<comparison predicate>”, are applied to:

RV1 <comp op> *RV2*

Access Rules

None.

General Rules

- 1) Let *R* be the result of the <row value predicand> and let *T* be the result of the <table subquery>.
- 2) The result of “*R* <comp op> <quantifier> *T*” is derived by the application of the implied <comparison predicate> “*R* <comp op> *RT*” to every row *RT* in *T*.

Case:

- a) If *T* is empty or if the implied <comparison predicate> is True for every row *RT* in *T*, then “*R* <comp op> <all> *T*” is True.
- b) If the implied <comparison predicate> is False for at least one row *RT* in *T*, then “*R* <comp op> <all> *T*” is False.

- c) If the implied <comparison predicate> is True for at least one row RT in T , then “ R <comp op> <some> T ” is True.
- d) If T is empty or if the implied <comparison predicate> is False for every row RT in T , then “ R <comp op> <some> T ” is False.
- e) If “ R <comp op> <quantifier> T ” is neither True nor False, then it is Unknown.

Conformance Rules

None.

NOTE 220 — If <equals operator> or <not equals operator> is specified, then the <quantified comparison predicate> is an equality operator and the Conformance Rules of [Subclause 9.9, “Equality operations”](#), apply. Otherwise, the <quantified comparison predicate> is an ordering operation, and the Conformance Rules of [Subclause 9.12, “Ordering operations”](#), apply.

8.10 <exists predicate>

Function

Specify a test for a non-empty set.

Format

```
<exists predicate> ::=  
  EXISTS <table subquery>
```

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let T be the result of the <table subquery>.
- 2) If the cardinality of T is greater than 0 (zero), then the result of the <exists predicate> is True; otherwise, the result of the <exists predicate> is False.

Conformance Rules

- 1) Without Feature T501, “Enhanced EXISTS predicate”, conforming SQL language shall not contain an <exists predicate> that simply contains a <table subquery> in which the <select list> of a <query specification> directly contained in the <table subquery> does not comprise either an <asterisk> or a single <derived column>.

8.11 <unique predicate>

Function

Specify a test for the absence of duplicate rows.

Format

```
<unique predicate> ::=  
    UNIQUE <table subquery>
```

Syntax Rules

- 1) Each column of user-defined type in the result of the <table subquery> shall have a comparison type.
- 2) Each column of the <table subquery> is an operand of a grouping operation. The Syntax Rules of [Subclause 9.10, “Grouping operations”](#), apply.

Access Rules

None.

General Rules

- 1) Let *T* be the result of the <table subquery>.
- 2) If there are no two rows in *T* such that the value of each column in one row is non-null and is not distinct from the value of the corresponding column in the other row, then the result of the <unique predicate> is True; otherwise, the result of the <unique predicate> is False.

Conformance Rules

- 1) Without Feature F291, “UNIQUE predicate”, conforming SQL language shall not contain a <unique predicate>.

NOTE 221 — The Conformance Rules of [Subclause 9.10, “Grouping operations”](#), also apply.

8.12 <normalized predicate>

Function

Determine whether a character string value is normalized.

Format

```
<normalized predicate> ::=
    <row value predicand> <normalized predicate part 2>

<normalized predicate part 2> ::=
    IS [ NOT ] [ <normal form> ] NORMALIZED
```

Syntax Rules

- 1) The <row value predicand> shall be a <row value constructor predicand> that is a <common value expression> *CVE*. The declared type of *CVE* shall be character string and the character set of *CVE* shall be UTF8, UTF16, or UTF32.
- 2) Case:
 - a) If <normal form> is specified, then let *NF* be <normal form>.
 - b) Otherwise, let *NF* be NFC.
- 3) The expression

CVE IS NOT *NF* NORMALIZED

is equivalent to

NOT (*CVE* IS *NF* NORMALIZED)

Access Rules

None.

General Rules

- 1) The result of *CVE* IS *NF* NORMALIZED is

Case:

 - a) If the value of *CVE* is the null value, then Unknown.
 - b) If the value of *CVE* is in the normalization form specified by *NF*, as defined by [Unicode15], then True.
 - c) Otherwise, False.

Conformance Rules

- 1) Without Feature T061, “UCS support”, conforming SQL language shall not contain a <normalized predicate>.
- 2) Without Feature F394, “Optional normal form specification”, conforming SQL language shall not contain <normal form>.

8.13 <match predicate>

Function

Specify a test for matching rows.

Format

```
<match predicate> ::=
  <row value predicand> <match predicate part 2>

<match predicate part 2> ::=
  MATCH [ UNIQUE ] [ SIMPLE | PARTIAL | FULL ] <table subquery>
```

Syntax Rules

- 1) The row type of the <row value predicand> and the row type of the <table subquery> shall be comparable.
- 2) Each field of <row value predicand> and each column of <table subquery> is an operand of an equality operation. The Syntax Rules of [Subclause 9.9, “Equality operations”](#), apply.
- 3) If neither SIMPLE, PARTIAL, nor FULL is specified, then SIMPLE is implicit.

Access Rules

None.

General Rules

- 1) Let R be the <row value predicand>.
- 2) If SIMPLE is specified or implicit, then

Case:

 - a) If R is the null value, then the <match predicate> is True.
 - b) Otherwise:
 - i) If the value of some field in R is the null value, then the <match predicate> is True.
 - ii) If the value of no field in R is the null value, then

Case:

 - 1) If UNIQUE is not specified and there exists a row RT_i of the <table subquery> such that

$$R = RT_i$$
 then the <match predicate> is True.

- 2) If UNIQUE is specified and there exists exactly one row RT_i in the result of evaluating the <table subquery> such that

$$R = RT_i$$

then the <match predicate> is True.

- 3) Otherwise, the <match predicate> is False.

- 3) If PARTIAL is specified, then

Case:

- a) If R is the null value, then the <match predicate> is True.
b) Otherwise,

Case:

- i) If the value of every field in R is the null value, then the <match predicate> is True.
ii) Otherwise,

Case:

- 1) If UNIQUE is not specified and there exists a row RT_i of the <table subquery> such that each non-null value of R equals its corresponding value in RT_i , then the <match predicate> is True.
2) If UNIQUE is specified and there exists exactly one row RT_i in the result of evaluating the <table subquery> such that each non-null value of R equals its corresponding value in RT_i , then the <match predicate> is True.
3) Otherwise, the <match predicate> is False.

- 4) If FULL is specified, then

Case:

- a) If R is the null value, then the <match predicate> is True.
b) Otherwise,

Case:

- i) If the value of every field in R is the null value, then the <match predicate> is True.
ii) If the value of no field in R is the null value, then

Case:

- 1) If UNIQUE is not specified and there exists a row RT_i of the <table subquery> such that

$$R = RT_i$$

then the <match predicate> is True.

8.13 <match predicate>

- 2) If UNIQUE is specified and there exists exactly one row RT_i in the result of evaluating the <table subquery> such that
$$R = RT_i$$
then the <match predicate> is True.
 - 3) Otherwise, the <match predicate> is False.
- iii) Otherwise, the <match predicate> is False.

Conformance Rules

- 1) Without Feature F741, “Referential MATCH types”, conforming SQL language shall not contain a <match predicate>.

NOTE 222 — The Conformance Rules of [Subclause 9.9](#), “Equality operations”, also apply.

8.14 <overlaps predicate>

Function

Specify a test for an overlap between two datetime periods.

Format

```
<overlaps predicate> ::=  
    <overlaps predicate part 1> <overlaps predicate part 2>  
  
<overlaps predicate part 1> ::=  
    <row value predicand 1>  
  
<overlaps predicate part 2> ::=  
    OVERLAPS <row value predicand 2>  
  
<row value predicand 1> ::=  
    <row value predicand>  
  
<row value predicand 2> ::=  
    <row value predicand>
```

Syntax Rules

- 1) The degrees of <row value predicand 1> and <row value predicand 2> shall both be 2.
- 2) The declared types of the first field of <row value predicand 1> and the first field of <row value predicand 2> shall both be datetime data types and these data types shall be comparable.

NOTE 223 — Two datetimes are comparable only if they have the same <primary datetime field>s; see [Subclause 4.6.2](#), “Datetimes”.

- 3) The declared type of the second field of each <row value predicand> shall be a datetime data type or INTERVAL.

Case:

- a) If the declared type is INTERVAL, then the precision of the declared type shall be such that the interval can be added to the datetime data type of the first column of the <row value predicand>.
- b) If the declared type is a datetime data type, then it shall be comparable with the datetime data type of the first column of the <row value predicand>.

Access Rules

None.

General Rules

- 1) If the value of <row value predicand 1> is the null value or the value of <row value predicand 2> is the null value, then the result of the <overlaps predicate> is Unknown and no further General Rules of this Subclause are applied.
- 2) Let $D1$ be the value of the first field of <row value predicand 1> and $D2$ be the value of the first field of <row value predicand 2>.
- 3) Case:
 - a) If the most specific type of the second field of <row value predicand 1> is a datetime data type, then let $E1$ be the value of the second field of <row value predicand 1>.
 - b) If the most specific type of the second field of <row value predicand 1> is INTERVAL, then let $I1$ be the value of the second field of <row value predicand 1>. Let $E1 = D1 + I1$.
- 4) If $D1$ is the null value or if $E1 < D1$, then let $S1 = E1$ and let $T1 = D1$. Otherwise, let $S1 = D1$ and let $T1 = E1$.
- 5) Case:
 - a) If the most specific type of the second field of <row value predicand 2> is a datetime data type, then let $E2$ be the value of the second field of <row value predicand 2>.
 - b) If the most specific type of the second field of <row value predicand 2> is INTERVAL, then let $I2$ be the value of the second field of <row value predicand 2>. Let $E2 = D2 + I2$.
- 6) If $D2$ is the null value or if $E2 < D2$, then let $S2 = E2$ and let $T2 = D2$. Otherwise, let $S2 = D2$ and let $T2 = E2$.
- 7) The result of the <overlaps predicate> is the result of the following expression:

```
( S1 > S2 AND NOT ( S1 >= T2 AND T1 >= T2 ) )  
OR  
( S2 > S1 AND NOT ( S2 >= T1 AND T2 >= T1 ) )  
OR  
( S1 = S2 AND ( T1 <> T2 OR T1 = T2 ) )
```

Conformance Rules

- 1) Without Feature F053, “OVERLAPS predicate”, conforming SQL language shall not contain an <overlaps predicate>.

8.15 <distinct predicate>

Function

Specify a test of whether two row values are distinct

Format

```
<distinct predicate> ::=
  <row value predicand 3> <distinct predicate part 2>

<distinct predicate part 2> ::=
  IS [ NOT ] DISTINCT FROM <row value predicand 4>

<row value predicand 3> ::=
  <row value predicand>

<row value predicand 4> ::=
  <row value predicand>
```

Syntax Rules

- 1) The two <row value predicand>s shall be of the same degree.
- 2) Let *respective values* be values with the same ordinal position.
- 3) The declared types of the respective values of the two <row value predicand>s shall be comparable.
- 4) Let *X* be the first <row value predicand> and let *Y* be the second <row value predicand>.
- 5) Each field of each <row value predicand> is an operand of an equality operation. The Syntax Rules of [Subclause 9.9, “Equality operations”](#), apply.
- 6) If <distinct predicate part 2> immediately contains NOT, then the <distinct predicate> is equivalent to:

NOT (*X* IS DISTINCT FROM *Y*)

Access Rules

None.

General Rules

- 1) Let *V1* be the value of <row value predicand 3> and let *V2* be the value of <row value predicand 4>.
Case:
 - a) If *V1* is the null value and *V2* is the null value, then the result is *False*.
 - b) If *V1* is the null value and *V2* is not the null value, or if *V1* is not the null value and *V2* is the null value, then the result is *True*.

c) Otherwise:

i) If *V1* and *V2* are values of a predefined type or a reference type, then

Case:

1) If *V1* and *V2* are not equal, then the result is True.

2) Otherwise, the result is False.

ii) If *V1* and *V2* are values of a user-defined type whose comparison form is RELATIVE or MAP, then

Case:

1) If the result of comparing *V1* and *V2* for equality according to Subclause 8.2, “<comparison predicate>”, is Unknown, then it is implementation-dependent whether the result is True or False.

2) If *V1* is not equal to *V2*, then the the result is True.

3) Otherwise the result is False.

iii) If *V1* and *V2* are values of a user-defined type whose comparison form is STATE, then

Case:

1) If the most specific types of *V1* and *V2* are different, then the result is True.

2) If there is an attribute *A* of their common most specific type such that the value of *A* in *V1* and the value of *A* in *V2* are distinct, then the result is True.

3) Otherwise the result is False.

iv) If *V1* and *V2* are values of row type, then

1) If at least one of their pairs of respective fields is distinct, then the result is True.

2) Otherwise, the result is False.

v) If *V1* and *V2* are of values of array type, then

Case:

1) If *V1* and *V2* do not have the same cardinality, then the result is True.

2) If *V1* and *V2* have the same cardinality and there exists at least one ordinal position *P* such that the array element at position *P* in *V1* is distinct from the array element at position *P* in *V2*, then the result is True.

3) Otherwise, the result is False.

vi) If *V1* and *V2* are values of multiset type, then

Case:

1) If there exists a value *V* in the element type of *V1* or *V2*, including the null value, such that the number of elements in *V1* that are not distinct from *V* does not equal the number of elements in *V2* that are not distinct from *V*, then the result is True.

- 2) Otherwise, the result is *False*.

NOTE 224 — “distinct” is defined in Subclause 3.1.6, “Definitions provided in Part 2”, and Subclause 4.1.5, “Properties of distinct”.

- 2) If two <row value predicand>s are not distinct, then they are said to be *duplicates*. If a number of <row value predicand>s are all duplicates of each other, then all except one are said to be *redundant duplicates*.

Conformance Rules

- 1) Without Feature T151, “DISTINCT predicate”, conforming SQL language shall not contain a <distinct predicate>.

NOTE 225 — The Conformance Rules of Subclause 9.9, “Equality operations”, also apply.

- 2) Without Feature T152, “DISTINCT predicate with negation”, conforming SQL language shall not contain a <distinct predicate part 2> that immediately contains NOT.

8.16 <member predicate>

Function

Specify a test of whether a value is a member of a multiset.

Format

```
<member predicate> ::=  
  <row value predicand> <member predicate part 2>  
  
<member predicate part 2> ::=  
  [ NOT ] MEMBER [ OF ] <multiset value expression>
```

Syntax Rules

- 1) Let *MVE* be the <multiset value expression> and let *ET* be the declared element type of *MVE*.
- 2) Case:
 - a) If the <row value predicand> is a <row value constructor predicand> that is a single <common value expression> or <boolean value expression> *CVE*, then let *X* be *CVE*.
 - b) Otherwise, let *X* be the <row value predicand>.
- 3) The declared type of *X* shall be comparable to *ET*.
- 4) *X* is an operand of an equality operation. The Syntax Rules of [Subclause 9.9, “Equality operations”](#), apply.
- 5) If <member predicate part 2> immediately contains NOT, then the <member predicate> is equivalent to

NOT (*X* MEMBER OF *MVE*)

Access Rules

None.

General Rules

- 1) Let *XV* be the value of *X*, and let *MV* be the value of *MVE*.
- 2) Let *N* be the result of `CARDINALITY (MVE)`.
- 3) The <member predicate>

XV MEMBER OF *MVE*

is evaluated as follows.

Case:

- a) If *N* is 0 (zero), then the <member predicate> is *False*.

- b) If either XV or MV is the null value, then the <member predicate> is Unknown.
- c) Otherwise, let ME_i for $1 \text{ (one)} \leq i \leq N$ be an enumeration of the elements of MV .

Case:

- i) If $CV = ME_i$ for some i , then the <member predicate> is True.
- ii) If ME_i is the null value for some i , then the <member predicate> is Unknown.
- iii) Otherwise, the <member predicate> is False.

Conformance Rules

- 1) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <member predicate>.

NOTE 226 — The Conformance Rules of Subclause 9.9, “Equality operations”, also apply.

8.17 <submultiset predicate>

Function

Specify a test of whether a multiset is a submultiset of another multiset.

Format

```
<submultiset predicate> ::=  
  <row value predicand> <submultiset predicate part 2>  
  
<submultiset predicate part 2> ::=  
  [ NOT ] SUBMULTISET [ OF ] <multiset value expression>
```

Syntax Rules

- 1) The <row value predicand> shall be a <row value constructor predicand> that is a single <common value expression> *CVE*. The declared type of *CVE* shall be a multiset type. Let *CVET* be the declared element type of *CVE*.
- 2) Let *MVE* be the <multiset value expression>. Let *MVET* be the declared element type of *MVE*.
- 3) *CVET* shall be comparable to *MVET*.
- 4) *CVE* and *MVE* are multiset operands of a multiset element grouping operation. The Syntax Rules of [Subclause 9.11, “Multiset element grouping operations”](#), apply.
- 5) If <submultiset predicate part 2> immediately contains NOT, then the <submultiset predicate> is equivalent to

NOT (*CVE* SUBMULTISET OF *MVE*)

Access Rules

None.

General Rules

- 1) Let *CV* be the value of *CVE*, and let *MV* be the value of *MVE*.
- 2) Let *M* be the result of CARDINALITY (*CV*), and let *N* be the result of CARDINALITY (*MV*).
- 3) The <submultiset predicate>

CVE SUBMULTISET OF *MVE*

is evaluated as follows.

Case:

- a) If *M* is 0 (zero), then the <submultiset predicate> is True.

b) If either CV or MV is the null value, then the <submultiset predicate> is Unknown.

c) Otherwise,

Case:

i) If $M > N$, then the <submultiset predicate> is False.

ii) If there exist an enumeration CE_i for $1 \text{ (one)} \leq i \leq M$ of the elements of CV and an enumeration ME_j for $1 \text{ (one)} \leq j \leq N$ of the elements of MV such that for all i , $1 \text{ (one)} \leq i \leq M$, $CE_i = ME_i$, then the <submultiset predicate> is True.

iii) If there exist an enumeration CE_i for $1 \text{ (one)} \leq i \leq M$ of the elements of CV and an enumeration ME_i for $1 \text{ (one)} \leq i \leq N$ of the elements of MV such that for all i , $1 \text{ (one)} \leq i \leq M$, $CE_i = ME_i$ is either True or Unknown, then the <submultiset predicate> is Unknown.

iv) Otherwise, the <submultiset predicate> is False.

Conformance Rules

- 1) Without Feature S275, “Advanced multiset support”, conforming SQL language shall not contain a <submultiset predicate>.

NOTE 227 — The Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

8.18 <set predicate>

Function

Specify a test of whether a multiset is a set (that is, does not contain any duplicates).

Format

```
<set predicate> ::=
  <row value predicand> <set predicate part 2>

<set predicate part 2> ::=
  IS [ NOT ] A SET
```

Syntax Rules

- 1) The <row value predicand> shall be a <row value constructor predicand> that is a single <common value expression> *CVE*. The declared type of *CVE* shall be a multiset type. Let *CVET* be the element type of *CVE*.
- 2) *CVE* is an operand of a multiset element grouping operation. The Syntax Rules of [Subclause 9.11](#), “Multiset element grouping operations”, apply.
- 3) If <set predicate part 2> immediately contains NOT, then the <set predicate> is equivalent to

$$\text{NOT (CVE IS A SET)}$$
- 4) If <set predicate part 2> does not immediately contain NOT, then the <set predicate> is equivalent to

$$\text{CARDINALITY (CVE) = CARDINALITY (SET (CVE))}$$

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <set predicate>.

NOTE 228 — The Conformance Rules of [Subclause 9.11](#), “Multiset element grouping operations”, also apply.

8.19 <type predicate>

Function

Specify a type test.

Format

```

<type predicate> ::=
    <row value predicand> <type predicate part 2>

<type predicate part 2> ::=
    IS [ NOT ] OF <left paren> <type list> <right paren>

<type list> ::=
    <user-defined type specification>
    [ { <comma> <user-defined type specification> }... ]

<user-defined type specification> ::=
    <inclusive user-defined type specification>
    | <exclusive user-defined type specification>

<inclusive user-defined type specification> ::=
    <path-resolved user-defined type name>

<exclusive user-defined type specification> ::=
    ONLY <path-resolved user-defined type name>

```

Syntax Rules

- 1) The <row value predicand> immediately contained in <type predicate> shall be a <row value constructor predicand> that is a <common value expression> *CVE*.
- 2) The declared type of *CVE* shall be a user-defined type.
- 3) For each <user-defined type name> *UDTN* contained in a <user-defined type specification>, the schema identified by the implicit or explicit schema name of *UDTN* shall include a user-defined type descriptor whose name is equivalent to the <qualified identifier> of *UDTN*.
- 4) Let the term *specified type* refer to a user-defined type that is specified by a <user-defined type name> contained in a <user-defined type specification>. A type specified by an <inclusive user-defined type specification> is *inclusively specified*; a type specified by an <exclusive user-defined type specification> is *exclusively specified*.
- 5) Let *T* be the type specified by <inclusive user-defined type specification> or <exclusive user-defined type specification>. *T* shall be a subtype of the declared type of *CVE*.

NOTE 229 — The term “subtype family” is defined in [Subclause 4.7.5](#), “Subtypes and supertypes”. If *T1* is a member of the subtype family of *T2*, then it follows that the subtype family of *T1* and the subtype family of *T2* are the same set of types.
- 6) Let *TL* be the <type list>.
- 7) A <type predicate> of the form

8.19 <type predicate>

CVE IS NOT
OF (*TL*)

is equivalent to

NOT (*CVE* IS
OF (*TL*))

Access Rules

None.

General Rules

- 1) Let *V* be the result of evaluating the <row value predicand>.
- 2) Let *ST* be the set consisting of every type that is either some exclusively specified type, or a subtype of some inclusively specified type.
- 3) Let *TPR* be the result of evaluating the <type predicate>.

Case:

- a) If *V* is the null value, then *TPR* is Unknown.
- b) If the most specific type of *V* is a member of *ST*, then *TPR* is True.
- c) Otherwise, *TPR* is False.

Conformance Rules

- 1) Without Feature S151, “Type predicate”, conforming SQL language shall not contain a <type predicate>.

8.20 <search condition>

Function

Specify a condition that is *True*, *False*, or *Unknown*, depending on the value of a <boolean value expression>.

Format

```
<search condition> ::=  
  <boolean value expression>
```

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) The result of the <search condition> is the result of the <boolean value expression>.

Conformance Rules

None.

(Blank page)

9 Additional common rules

9.1 Retrieval assignment

This Subclause is modified by Subclause 9.1, “Retrieval assignment”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.1, “Retrieval assignment”, in ISO/IEC 9075-14.

Subclause Signature

```
"Retrieval assignment" [Syntax Rules] (  
Parameter: "TARGET" ,  
Parameter: "VALUE"  
)
```

Function

Specify rules for assignments to targets that do not support null values or that support null values with indicator parameters (*e.g.*, assigning SQL-data to host parameters or host variables).

Syntax Rules

- 1) Let *T* be the *TARGET* and let *V* be the *VALUE* in an application of this Subclause. Let the declared types of *T* and *V* be *TD* and *SD*, respectively.
- 2) If *TD* is numeric, boolean, datetime, interval, or a user-defined type, then either *SD* shall be assignable to *TD* or there shall exist an appropriate user-defined cast function *UDCF* from *SD* to *TD*.

NOTE 230 — “Appropriate user-defined cast function” is defined in Subclause 4.11, “Data conversions”.

- 3) If *TD* is binary string, then

Case:

- a) If *T* is either a locator parameter of an external routine, a locator variable, or a host parameter that is a binary large object locator parameter, then *SD* shall be BINARY LARGE OBJECT and *SD* shall be assignable to *TD*.
- b) Otherwise, either *SD* shall be assignable to *TD* or there shall exist an appropriate user-defined cast function *UDCF* from *SD* to *TD*.

- 4) If *TD* is character string, then

Case:

9.1 Retrieval assignment

- a) If T is either a locator parameter of an external routine, a locator variable, or a host parameter that is a character large object locator parameter, then SD shall be CHARACTER LARGE OBJECT and SD shall be assignable to TD .
 - b) Otherwise, either SD shall be assignable to TD or there shall exist an appropriate user-defined cast function $UDCF$ from SD to TD .
- 5) If the declared type of T is a reference type, then the declared type of V shall be a reference type whose referenced type is a subtype of the referenced type of T .
- 6) If the declared type of T is a row type, then:
- a) The declared type of V shall be a row type.
 - b) The degree of V shall be the same as the degree of T . Let n be that degree.
 - c) Let TT_i , $1 \text{ (one)} \leq i \leq n$, be the declared type of the i -th field of T , let VT_i be the declared type of the i -th field of V , let TI_i be an arbitrary target whose declared type is TT_i , and let VI_i be an arbitrary expression whose declared type is VT_i . For each i , $1 \text{ (one)} \leq i \leq n$, the Syntax Rules of this Subclause are applied with TI_i as *TARGET* and VI_i as *VALUE*.
- 7) If the declared type of T is a collection type, then:
- a) If the declared type of T is an array type, then the declared type of V shall be an array type.
 - b) If the declared type of T is a multiset type, then the declared type of V shall be a multiset type.
 - c) Let TT be the element type of the declared type of T , let VT be the element type of the declared type of V , let TI be an arbitrary target whose declared type is TT , and let VI be an arbitrary expression whose declared type is VT . The Syntax Rules of this Subclause are applied with TI as *TARGET* and VI as *VALUE*.

Access Rules

None.

General Rules

- 1) If the declared type of V is not assignable to the declared type of T , then for the remaining General Rules of this Subclause V is effectively replaced by the result of evaluating the expression $UDCF(V)$.
- 2) If V is the null value and T is a host parameter, then
Case:
 - a) If an indicator parameter is specified for T , then that indicator parameter is set to -1 .
 - b) If no indicator parameter is specified for T , then an exception condition is raised: *data exception — null value, no indicator parameter*.
- 3) If V is the null value and T is a host variable, then
Case:
 - a) If an indicator variable is specified for T , then that indicator variable is set to -1 .

- b) If no indicator variable is specified for T , then an exception condition is raised: *data exception — null value, no indicator parameter*.

- 4) If V is not the null value, T is a host parameter, and T has an indicator parameter, then

Case:

- a) If the declared type of T is character string or binary string and the length M in characters or octets, respectively, of V is greater than the length in characters or octets, respectively, of T , then the indicator parameter is set to M . If M exceeds the maximum value that the indicator parameter can contain, then an exception condition is raised: *data exception — indicator overflow*.
- b) Otherwise, the indicator parameter is set to 0 (zero).

- 5) If V is not the null value, T is a host variable, and T has an indicator variable, then

Case:

- a) If the declared type of T is character string or binary string and the length in characters or octets, respectively, M of V is greater than the length in characters or octets, respectively, of T , then the indicator variable is set to M . If M exceeds the maximum value that the indicator variable can contain, then an exception condition is raised: *data exception — indicator overflow*.
- b) Otherwise, the indicator variable is set to 0 (zero).

- 6) 09 14 If V is not the null value, then

Case:

- a) If the declared type of T is fixed-length character string with length in characters L and the length in characters of V is equal to L , then the value of T is set to V .
- b) If the declared type of T is fixed-length character string with length in characters L , and the length in characters of V is greater than L , then the value of T is set to the first L characters of V and a completion condition is raised: *warning — string data, right truncation*.
- c) If the declared type of T is fixed-length character string with length in characters L , and the length in characters M of V is smaller than L , then the first M characters of T are set to V , and the last $L-M$ characters of T are set to <space>s.
- d) If the declared type of T is variable-length character string and the length in characters M of V is not greater than the maximum length in characters of T , then the value of T is set to V and the length in characters of T is set to M .
- e) If the declared type of T is variable-length character string and the length in characters of V is greater than the maximum length in characters L of T , then the value of T is set to the first L characters of V , then the length in characters of T becomes L , and a completion condition is raised: *warning — string data, right truncation*.
- f) If the declared type of T is a character large object type and the length in characters M of V is not greater than the maximum length in characters of T , then the value of T is set to V and the length in characters of T is set to M .
- g) If the declared type of T is a character large object type and the length in characters of V is greater than the maximum length in characters L of T , then the value of T is set to the first L characters of V , the length in characters of T becomes L , and a completion condition is raised: *warning — string data, right truncation*.

9.1 Retrieval assignment

- h) If the declared type of *T* is fixed-length binary string with length in octets *L* and the length in octets of *V* is equal to *L*, then the value of *T* is set to *V*.
- i) If the declared type of *T* is fixed-length binary string with length in octets *L*, and the length in octets of *V* is greater than *L*, then the value of *T* is set to the first *L* octets of *V* and a completion condition is raised: *warning — string data, right truncation*.
- j) If the declared type of *T* is fixed-length binary string with length in octets *L*, and the length in octets *M* of *V* is less than *L*, then the first *M* octets of *T* are set to *V*, and the last *L–M* octets of *T* are set to X'00's.
- k) If the declared type of *T* is variable-length binary string and the length in octets *M* of *V* is not greater than the maximum length in octets of *T*, then the value of *T* is set to *V* and the length in octets of *T* is set to *M*.
- l) If the declared type of *T* is variable-length binary string and the length in octets of *V* is greater than the maximum length in octets *L* of *T*, then the value of *T* is set to the first *L* octets of *V*, then the length in octets of *T* becomes *L*, and a completion condition is raised: *warning — string data, right truncation*.
- m) If the declared type of *T* is binary large object string and the length in octets *M* of *V* is not greater than the maximum length in octets of *T*, then the value of *T* is set to *V* and the length in octets of *T* is set to *M*.
- n) If the declared type of *T* is binary large object string and the length in octets of *V* is greater than the maximum length in octets *L* of *T*, then the value of *T* is set to the first *L* octets of *V*, the length in octets of *T* becomes *L*, and a completion condition is raised: *warning — string data, right truncation*.
- o) If the declared type of *T* is numeric, then

Case:

- i) If *V* is a member of the declared type of *T*, then *T* is set to *V*.
 - ii) If a member of the declared type of *T* can be obtained from *V* by rounding or truncation, then *T* is set to that value. If the declared type of *T* is exact numeric, then it is implementation-defined whether the approximation is obtained by rounding or by truncation.
 - iii) Otherwise, an exception condition is raised: *data exception — numeric value out of range*.
- p) If the declared type of *T* is boolean, then the value of *T* is set to *V*.
- q) If the declared type *DT* of *T* is datetime, then:
- i) If only one of *DT* and the declared type of *V* is datetime with time zone, then *V* is effectively replaced by

CAST (*V* AS *DT*)
 - ii) Case:
 - 1) If *V* is a member of the declared type of *T*, then *T* is set to *V*.
 - 2) If a member of the declared type of *T* can be obtained from *V* by rounding or truncation, then *T* is set to that value. It is implementation-defined whether the approximation is obtained by rounding or truncation.
 - 3) Otherwise, an exception condition is raised: *data exception — datetime field overflow*.

- r) If the declared type of T is interval, then

Case:

- i) If V is a member of the declared type of T , then T is set to V .
- ii) If a member of the declared type of T can be obtained from V by rounding or truncation, then T is set to that value. It is implementation-defined whether the approximation is obtained by rounding or by truncation.
- iii) Otherwise, an exception condition is raised: *data exception — interval field overflow*.

- s) If the declared type of T is a row type, then:

- i) Let n be the degree of T .
- ii) For i ranging from 1 (one) to n , the General Rules of this Subclause are applied to the i -th element of T and the i -th element of V as *TARGET* and *VALUE*, respectively.

- t) If the declared type of T is a user-defined type, then the value of T is set to V .

- u) If the declared type of T is a reference type, then the value of T is set to V .

- v) If the declared type of T is an array type, then

Case:

- i) If the maximum cardinality L of T is equal to the cardinality M of V , then the elements of T are set to the values of the corresponding elements of V by applying the General Rules of this Subclause to each pair of elements with the element of T as *TARGET* and the element of V as *VALUE*.
- ii) If the maximum cardinality L of T is smaller than the cardinality M of V , then the elements of T are set to the values of the first L corresponding elements of V by applying the General Rules of this Subclause to each pair of elements with the element of T as *TARGET* and the element of V as *VALUE*; a completion condition is raised: *warning — array data, right truncation*.
- iii) If the maximum cardinality L of T is greater than the cardinality M of V , then the M first elements of T are set to the values of the corresponding elements of V by applying the General Rules of this Subclause to each pair of elements with the element of T as *TARGET* and the element of V as *VALUE*. The cardinality of the value of T is M .

NOTE 231 — The maximum cardinality L of T is unchanged.

- w) If the declared type of T is a multiset type, then the value of T is set to V .

Conformance Rules

None.

9.2 Store assignment

This Subclause is modified by Subclause 9.2, “Store assignment”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.2, “Store assignment”, in ISO/IEC 9075-14.

Subclause Signature

```
"Store assignment" [Syntax Rules] (
Parameter: "TARGET",
Parameter: "VALUE"
)
```

Function

Specify rules for assignments where the target permits null without the use of indicator parameters or indicator variables, such as storing SQL-data or setting the value of SQL parameters.

Syntax Rules

- 1) Let T be the *TARGET* and let V be the *VALUE* in an application of this Subclause. Let the declared types of T and V be TD and SD , respectively.
- 2) If TD is character string, binary string, numeric, boolean, datetime, interval, or a user-defined type, then either SD shall be assignable to TD or there shall exist an appropriate user-defined cast function $UDCF$ from SD to TD .

NOTE 232 — “Appropriate user-defined cast function” is defined in Subclause 4.11, “Data conversions”.

- 3) If the declared type of T is a reference type, then the declared type of V shall be a reference type whose referenced type is a subtype of the referenced type of T .
- 4) If the declared type of T is a row type, then:
 - a) The declared type of V shall be a row type.
 - b) The degree of V shall be the same as the degree of T . Let n be that degree.
 - c) Let TT_i , $1 \text{ (one)} \leq i \leq n$, be the declared type of the i -th field of T , let VT_i be the declared type of the i -th field of V , let TI_i be an arbitrary target whose declared type is TT_i , and let VI_i be an arbitrary expression whose declared type is VT_i . For each i , $1 \text{ (one)} \leq i \leq n$, the Syntax Rules of this Subclause apply to TI_i VI_i , as *TARGET* and *VALUE*, respectively.
- 5) If the declared type of T is a collection type, then:
 - a) If the declared type of T is an array type, then the declared type of V shall be an array type.
 - b) If the declared type of T is a multiset type, then the declared type of V shall be a multiset type.
 - c) Let TT be the element type of the declared type of T , let VT be the element type of the declared type of V , let TI be an arbitrary target whose declared type is TT , and let VI be an arbitrary expression whose declared type is VT . The Syntax Rules of this Subclause apply to TI and VI , as *TARGET* and *VALUE*, respectively.

Access Rules

None.

General Rules

1) If the declared type of V is not assignable to the declared type of T , then for the remaining General Rules of this Subclause V is effectively replaced by the result of evaluating the expression $UDCF(V)$.

2) Case:

a) If V is the null value, then

Case:

i) If V is specified using NULL, then T is set to the null value.

ii) If V is a host parameter and contains an indicator parameter, then

Case:

1) If the value of the indicator parameter is equal to -1 , then T is set to the null value.

2) If the value of the indicator parameter is less than -1 , then an exception condition is raised:
data exception — invalid indicator parameter value.

iii) If V is a host variable and contains an indicator variable, then

Case:

1) If the value of the indicator variable is equal to -1 , then T is set to the null value.

2) If the value of the indicator variable is less than -1 , then an exception condition is raised:
data exception — invalid indicator parameter value.

iv) Otherwise, T is set to the null value.

b) 09 14 Otherwise,

Case:

i) If the declared type of T is fixed-length character string with length in characters L and the length in characters of V is equal to L , then the value of T is set to V .

ii) If the declared type of T is fixed-length character string with length in characters L and the length in characters M of V is larger than L , then

Case:

1) If the rightmost $M-L$ characters of V are all <space>s, then the value of T is set to the first L characters of V .

2) If one or more of the rightmost $M-L$ characters of V are not <space>s, then an exception condition is raised: *data exception — string data, right truncation.*

- iii) If the declared type of T is fixed-length character string with length in characters L and the length in characters M of V is less than L , then the first M characters of T are set to V and the last $L-M$ characters of T are set to <space>s.
- iv) If the declared type of T is variable-length character string and the length in characters M of V is not greater than the maximum length in characters of T , then the value of T is set to V and the length in characters of T is set to M .
- v) If the declared type of T is variable-length character string and the length in characters M of V is greater than the maximum length in characters L of T , then

Case:

- 1) If the rightmost $M-L$ characters of V are all <space>s, then the value of T is set to the first L characters of V and the length in characters of T is set to L .
- 2) If one or more of the rightmost $M-L$ characters of V are not <space>s, then an exception condition is raised: *data exception — string data, right truncation*.
- vi) If the declared type of T is a character large object type and the length in characters M of V is not greater than the maximum length in characters of T , then the value of T is set to V and the length in characters of T is set to M .
- vii) If the declared type of T is a character large object type and the length in characters M of V is greater than the maximum length in characters L of T , then

Case:

- 1) If the rightmost $M-L$ characters of V are all <space>s, then the value of T is set to the first L characters of V and the length in characters of T is set to L .
- 2) If one or more of the rightmost $M-L$ characters of V are not <space>s, then an exception condition is raised: *data exception — string data, right truncation*.
- viii) If the declared type of T is fixed-length binary string with length in octets L and the length in octets of V is equal to L , then the value of T is set to V .
- ix) If the declared type of T is fixed-length binary string with length in octets L and the length in octets M of V is larger than L , then

Case:

- 1) If the rightmost $M-L$ octets of V are all equal to X'00', then the value of T is set to the first L octets of V .
- 2) If one or more of the rightmost $M-L$ octets of V are not equal to X'00', then an exception condition is raised: *data exception — string data, right truncation*.
- x) If the declared type of T is fixed-length binary string with length in octets L and the length in octets M of V is less than L , then the first M octets of T are set to V and the last $L-M$ octets of T are set to X'00's.
- xi) If the declared type of T is variable-length binary string and the length in octets M of V is not greater than the maximum length in octets of T , then the value of T is set to V and the length in octets of T is set to M .

- xii) If the declared type of T is variable-length binary string and the length in octets M of V is greater than the maximum length in octets L of T , then

Case:

- 1) If the rightmost $M-L$ octets of V are all equal to X'00', then the value of T is set to the first L octets of V and the length in octets of T is set to L .
- 2) If one or more of the rightmost $M-L$ octets of V are not equal to X'00', then an exception condition is raised: *data exception — string data, right truncation*.

- xiii) If the declared type of T is binary large object string and the length in octets M of V is not greater than the maximum length in octets of T , then the value of T is set to V and the length in octets of T is set to M .

- xiv) If the declared type of T is binary large object string and the length in octets M of V is greater than the maximum length in octets L of T , then

Case:

- 1) If the rightmost $M-L$ octets of V are all equal to X'00', then the value of T is set to the first L octets of V and the length in octets of T is set to L .
- 2) If one or more of the rightmost $M-L$ octets of V are not equal to X'00', then an exception condition is raised: *data exception — string data, right truncation*.

- xv) If the declared type of T is numeric, then

Case:

- 1) If V is a member of the declared type of T , then T is set to V .
- 2) If a member of the declared type of T can be obtained from V by rounding or truncation, then T is set to that value. If the declared type of T is exact numeric, then it is implementation-defined whether the approximation is obtained by rounding or by truncation.
- 3) Otherwise, an exception condition is raised: *data exception — numeric value out of range*.

- xvi) If the declared type DT of T is datetime, then

- 1) If only one of DT and the declared type of V is datetime with time zone, then V is effectively replaced by

CAST (V AS DT)

- 2) Case:

- A) If V is a member of the declared type of T , then T is set to V .
- B) If a member of the declared type of T can be obtained from V by rounding or truncation, then T is set to that value. It is implementation-defined whether the approximation is obtained by rounding or truncation.
- C) Otherwise, an exception condition is raised: *data exception — datetime field overflow*.

- xvii) If the declared type of T is interval, then

Case:

- 1) If V is a member of the declared type of T , then T is set to V .
 - 2) If a member of the declared type of T can be obtained from V by rounding or truncation, then T is set to that value. It is implementation-defined whether the approximation is obtained by rounding or by truncation.
 - 3) Otherwise, an exception condition is raised: *data exception — interval field overflow*.
- xviii) If the declared type of T is boolean, then the value of T is set to V .
- xix) If the declared type of T is a row type, then:
- 1) Let n be the degree of T .
 - 2) For i ranging from 1 (one) to n , the General Rules of this Subclause are applied to the i -th element of T and the i -th element of V as *TARGET* and *VALUE*, respectively.
- xx) If the declared type of T is a user-defined type, then the value of T is set to V .
- xxi) If the declared type of T is a reference type, then the value of T is set to V .
- xxii) If the declared type of T is an array type, then

Case:

- 1) If the maximum cardinality L of T is equal to the cardinality M of V , then the elements of T are set to the values of the corresponding elements of V by applying the General Rules of this Subclause to each pair of elements with the element of T as *TARGET* and the element of V as *VALUE*.
- 2) If the maximum cardinality L of T is smaller than the cardinality M of V , then

Case:

- A) If the rightmost $M-L$ elements of V are all null, then the elements of T are set to the values of the first L corresponding elements of V by applying the General Rules of this Subclause to each pair of elements with the element of T as *TARGET* and the element of V as *VALUE*.
- B) If one or more of the rightmost $M-L$ elements of V are not the null value, then an exception condition is raised: *data exception — array data, right truncation*.
- 3) If the maximum cardinality L of T is greater than the cardinality M of V , then the M first elements of T are set to the values of the corresponding elements of V by applying the General Rules of this Subclause to each pair of elements with the element of T as *TARGET* and the element of V as *VALUE*. The cardinality of the value of T is set to M .

NOTE 233 — The maximum cardinality L of T is unchanged.

- xxiii) If the declared type of T is a multiset type, then the value of T is set to V .

Conformance Rules

None.

9.3 Result of data type combinations

*This Subclause is modified by Subclause 9.3, “Result of data type combinations”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 10.3, “Result of data type combinations”, in ISO/IEC 9075-14.*

Subclause Signature

“Result of data type combinations” [Syntax Rules] (
Parameter: “DTSET”
) Returns: “RESTYPE”

Function

Specify the result data type of the result of certain combinations of values of compatible data types, such as <case expression>s, <collection value expression>s, or a column in the result of a <query expression>.

Syntax Rules

- 1) Let *IDTS* be the *DTSET* in an application of this Subclause. The result data type that is the result of the application of this Subclause is *RESTYPE*.
- 2) Let *DTS* be the set of data types in *IDTS* excluding any data types that are undefined. If the cardinality of *DTS* is 0 (zero), then the result data type is undefined and no further Rules of this Subclause are evaluated.

NOTE 234 — The notion of “undefined data type” is defined in Subclause 20.6, “<prepare statement>”.

- 3) Case:
 - a) If any of the data types in *DTS* is character string, then:
 - i) All data types in *DTS* shall be character string, and all of them shall have the same character repertoire. The character set of the result is the character set of the data type in *DTS* that has the character encoding form with the highest precedence.
 - ii) The collation derivation and declared type collation of the result are determined as follows.
Case:
 - 1) If some data type in *DTS* has an *explicit* collation derivation and declared type collation *ECI*, then every data type in *DTS* that has an *explicit* collation derivation shall have a declared type collation that is *ECI*. The collation derivation is *explicit* and the collation is *ECI*.
 - 2) If every data type in *DTS* has an *implicit* collation derivation, then
Case:
 - A) If every data type in *DTS* has the same declared type collation *ICI*, then the collation derivation is *implicit* and the declared type collation is *ICI*.
 - B) Otherwise, the collation derivation is *none*.
 - 3) Otherwise, the collation derivation is *none*.

9.3 Result of data type combinations

iii) Case:

- 1) If any of the data types in *DTS* is a character large object type, then the result data type is a character large object type with maximum length in characters equal to the maximum of the lengths in characters and maximum lengths in characters of the data types in *DTS*.
- 2) If any of the data types in *DTS* is variable-length character string, then the result data type is variable-length character string with maximum length in characters equal to the maximum of the lengths in characters and maximum lengths in characters of the data types in *DTS*.
- 3) Otherwise, the result data type is fixed-length character string with length in characters equal to the maximum of the lengths in characters of the data types in *DTS*.

b) If any of the data types in *DTS* is binary string, then

Case:

- i) If any of the data types in *DTS* is a binary large object type, then the result data type is a binary large object type with maximum length in octets equal to the maximum of the lengths in octets and the maximum lengths in octets of the data types in *DTS*.
 - ii) If any of the data types in *DTS* is variable-length binary string, then the result data type is variable-length binary string with maximum length in octets equal to the maximum of the lengths in octets and the maximum lengths in octets of the data types in *DTS*.
 - iii) Otherwise, the result data type is fixed-length binary string with length in octets equal to the maximum of the lengths in octets of the data types in *DTS*.
- c) If all of the data types in *DTS* are exact numeric, then the result data type is exact numeric with implementation-defined precision and with scale equal to the maximum of the scales of the data types in *DTS*.
 - d) If any data type in *DTS* is approximate numeric, then each data type in *DTS* shall be numeric and the result data type is approximate numeric with implementation-defined precision.
 - e) If some data type in *DTS* is a datetime data type, then every data type in *DTS* shall be a datetime data type having the same datetime fields. The result data type is a datetime data type having the same datetime fields, whose fractional seconds precision is the largest of the fractional seconds precisions in *DTS*. If some data type in *DTS* has a time zone displacement value, then the result has a time zone displacement value; otherwise, the result does not have a time zone displacement value.
 - f) If any data type in *DTS* is interval, then each data type in *DTS* shall be interval. If the precision of any data type in *DTS* specifies YEAR or MONTH, then the precision of each data type shall specify only YEAR or MONTH. If the precision of any data type in *DTS* specifies DAY, HOUR, MINUTE, or SECOND(*N*), then the precision of no data type of *DTS* shall specify the <primary datetime field>s YEAR and MONTH. The result data type is interval with precision “*S TO E*”, where *S* and *E* are the most significant of the <start field>s and the least significant of the <end field>s of the data types in *DTS*, respectively.
 - g) 09 14 If any data type in *DTS* is boolean, then each data type in *DTS* shall be boolean. The result data type is boolean.
 - h) If any data type in *DTS* is a row type, then each data type in *DTS* shall be a row type with the same degree and the data type of each field in the same ordinal position of every row type shall be comparable. The result data type is a row type defined by an ordered sequence of (<field name>, data type) pairs FD_i , where data type is the data type resulting from the application of this Subclause to the set of data

types of fields in the same ordinal position as FD_i in every row type in DTS and <field name> is determined as follows.

Case:

- i) If the names of fields in the same ordinal position as FD_i in every row type in DTS is F , then the <field name> in FD_i is F .
 - ii) Otherwise, the <field name> in FD_i is implementation-dependent.
- i) If any data type in DTS is an array type then every data type in DTS shall be an array type. The data type of the result is array type with element data type ETR , where ETR is the data type resulting from the application of this Subclause to the set of element types of the array types of DTS , and maximum cardinality equal to the maximum of the maximum cardinalities of the data types in DTS .
 - j) If any data type in DTS is a multiset type then every data type in DTS shall be a multiset type. The data type of the result is multiset type with element data type ETR , where ETR is the data type resulting from the application of this Subclause to the set of element types of the multiset types of DTS .
 - k) If any data type in DTS is a reference type, then there shall exist a subtype family STF such that each data type in DTS is a member of STF . Let RT be the minimal common supertype of each data type in DTS .

Case:

- i) If the data type descriptor of every data type in DTS includes the name of a referenceable table identifying the scope of the reference type, and every such name is equivalent to some name STN , then result data type is:

$$RT \text{ SCOPE } (STN)$$
 - ii) Otherwise, the result data type is RT .
- l) Otherwise, there shall exist a subtype family STF such that each data type in DTS is a member of STF . The result data type is the minimal common supertype of each data type in DTS .

NOTE 235 — *Minimal common supertype* is defined in Subclause 4.7.5, “Subtypes and supertypes”.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.4 Subject routine determination

Subclause Signature

"Subject routine determination" [Syntax Rules] (
 Parameter: "SET OF SQL-INVOKED ROUTINES",
 Parameter: "SQL ARGUMENT LIST"
) Returns: "SET OF SUBJECT ROUTINES"

Function

Determine the subject routine of a given routine invocation.

Syntax Rules

- 1) Let *SR* be the *SET OF SQL-INVOKED ROUTINES* and let *AL* be the *SQL ARGUMENT LIST* in an application of this Subclause. The set of SQL-invoked routines that is the result of the application of this Subclause is *SET OF SUBJECT ROUTINES*.
- 2) Let *n* be the number of SQL-invoked routines in *SR*. Let R_i , 1 (one) $\leq i \leq n$, be the *i*-th SQL-invoked routine in *SR* in the ordering of *SR*.
- 3) Let *m* be the number of SQL arguments in *AL*. Let A_j , 1 (one) $\leq j \leq m$, be the *j*-th SQL argument in *AL*.
- 4) For A_j , 1 (one) $\leq j \leq m$, then let $SDTA_j$ be the declared type of A_j .
- 5) Let $SDTP_{i,j}$ be the type designator of the declared type of the *j*-th SQL parameter of R_i .
- 6) For *r* varying from 1 (one) to *m*, if A_r is not a <dynamic parameter specification> and if there is more than one SQL-invoked routine in *SR*, then for each pair of SQL-invoked routines $\{ R_p, R_q \}$ in *SR*:
 - a) The Syntax Rules of [Subclause 9.5, "Type precedence list determination"](#), are applied with $SDTA_r$ as *DATA TYPE*, yielding TPL_r returned from *TYPE PRECEDENCE LIST*.
 - b) If $SDTP_{p,r} \prec SDTP_{q,r}$ in TPL_r , then eliminate R_q from *SR*.
- 7) The set of SQL-invoked routines remaining in *SR* is the result of the application of this Subclause.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.5 Type precedence list determination

This Subclause is modified by Subclause 9.4, “Type precedence list determination”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.4, “Type precedence list determination”, in ISO/IEC 9075-14.

Subclause Signature

“Type precedence list determination” [Syntax Rules] (

Parameter: “DATA TYPE”

) Returns: “TYPE PRECEDENCE LIST”

Function

Determine the type precedence list of a given type.

Syntax Rules

- 1) Let *DT* be the *DATA TYPE* in an application of this Subclause. The *type precedence list TPL* that is the result of the application of this Subclause is *TYPE PRECEDENCE LIST*.
- 2) The list of *type designators* determined by the Syntax Rules of this Subclause is the result of the application of this Subclause.
- 3) Let “*A* < *B*” represent “*A* has precedence over *B*” and let “*A* \simeq *B*” represent “*A* has the same precedence as *B*”.
- 4) If *DT* is a user-defined type, then:
 - a) Let *ST* be the set of supertypes of *DT*. Let *n* be the number of data types in *ST*.
 - b) For any two data types *TA* and *TB* in *ST*, *TA* < *TB* if and only if *TA* is a proper subtype of *TB*.
 - c) Let *T*₁ be *DT* and let *T*_{*i*+1}, 1 (one) ≤ *i* ≤ *n*–1, be the direct supertype of *T*_{*i*}.
 - d) Let *DTN*_{*i*}, 1 (one) ≤ *i* ≤ *n*, be the data type designator of *T*_{*i*}.

NOTE 236 — The type designator of a user-defined type is the type name included in its user-defined type descriptor.

- e) *TPL* is *DTN*₁, *DTN*₂, ..., *DTN*_{*n*}.
- 5) If *DT* is fixed-length character string, then *TPL* is

CHARACTER, CHARACTER VARYING,
CHARACTER LARGE OBJECT
- 6) If *DT* is variable-length character string, then *TPL* is

CHARACTER VARYING, CHARACTER
LARGE OBJECT
- 7) If *DT* is fixed-length binary string, then *TPL* is

BINARY, BINARY VARYING, BINARY LARGE OBJECT

- 8) If DT is variable-length binary string, then TPL is

BINARY VARYING, BINARY LARGE OBJECT

- 9) If DT is numeric, then:

- a) Let NDT be the following set of numeric types: NUMERIC, DECIMAL, SMALLINT, INTEGER, BIGINT, REAL, FLOAT, and DOUBLE PRECISION. For each type T in NDT , the *effective binary precision* is defined as follows.

Case:

- i) If T is DECIMAL or NUMERIC, then the effective binary precision is the product of $\log_2(10)$ and the implementation-defined maximum precision of T .
 - ii) If T is FLOAT, then the effective binary precision is the implementation-defined maximum precision of T .
 - iii) If the radix of T is decimal, then the effective binary precision is the product of $\log_2(10)$ and the implementation-defined precision of T .
 - iv) Otherwise, the effective binary precision is the implementation-defined precision of T .
- b) Let PTC be the set of all precedence relationships determined as follows: For any two types $T1$ and $T2$, not necessarily distinct, in NDT ,

Case:

- i) If $T1$ is exact numeric and $T2$ is approximate numeric, then $T1 < T2$.
 - ii) If $T1$ is approximate numeric and $T2$ is exact numeric, then $T1 > T2$.
 - iii) If the effective binary precision of $T1$ is greater than the effective binary precision of $T2$, then $T2 < T1$.
 - iv) If the effective binary precision of $T1$ equals the effective binary precision of $T2$, then $T2 \simeq T1$.
 - v) Otherwise, $T1 < T2$.
- c) TPL is determined as follows:
- i) TPL is initially empty.
 - ii) Let ST be the set of types containing DT and every type T in NDT for which the precedence relationship $DT < T$ or $DT \simeq T$ is in PTC .
 - iii) Let n be the number of types in ST .
 - iv) For i ranging from 1 (one) to n :
 - 1) Let NT be the set of types T_k in ST such that there is no other type T_j in ST for which $T_j < T_k$ according to PTC .
 - 2) Case:

9.5 Type precedence list determination

- A) If there is exactly one type T_k in NT , then T_k is placed next in TPL and all relationships of the form $T_k < T_r$ are removed from PTC , where T_r is any type in ST .
- B) If there is more than one type T_k in NT , then every type T_s in NT is assigned the same position in TPL as T_k and all relationships of the forms $T_k < T_r$, $T_k \simeq T_r$, $T_s < T_r$, and $T_s \simeq T_r$ are removed from PTC , where T_r is any type in ST .

10) If DT specifies a year-month interval type, then TPL is

INTERVAL YEAR

11) If DT specifies a day-time interval type, then TPL is

INTERVAL DAY

12) If DT specifies DATE, then TPL is

DATE

13) If DT specifies TIME, then TPL is

TIME

14) If DT specifies TIMESTAMP, then TPL is

TIMESTAMP

15) If DT specifies BOOLEAN, then TPL is

BOOLEAN

16) If DT is a collection type, then let CTC be the kind of collection (either ARRAY or MULTISSET) specified in DT .

Let n be the number of elements in the type precedence list for the element type of DT . For i ranging from 1 (one) to n , let RIO_i be the i -th such element. TPL is

$RIO_1 \text{ } CTC,$
 $RIO_2 \text{ } CTC, \dots,$
 $RIO_n \text{ } CTC$

17) If DT is a reference type, then let n be the number of elements in the type precedence list for the referenced type of DT . For i ranging from 1 (one) to n , let KAW_i be the i -th such element. TPL is

$REF(KAW_1),$
 $REF(KAW_2), \dots,$
 $REF(KAW_n)$

18) If DT is a row type, then TPL is

ROW

NOTE 237 — This rule is placed only to avoid the confusion that might arise if row types were not mentioned in this Subclause. As a row type cannot be used as a <parameter type>, the type precedence list of a row type is never referenced.

Conformance Rules

None.

9.6 Host parameter mode determination

Subclause Signature

```
"Host parameter mode determination" [Syntax Rules] (
Parameter: "HOST PARAM DECL" ,
Parameter: "SQL PROC STMT"
)
```

Function

Determine the parameter mode for a given host parameter.

Syntax Rules

- 1) Let *PD* be the *HOST PARAM DECL* and let *SPS* be the *SQL PROC STMT* in an application of this Subclause.
- 2) Let *P* be the host parameter specified by *PD* and let *PN* be the <host parameter name> immediately contained in *PD*.
- 3) Whether *P* is an input host parameter, an output host parameter, or both an input host parameter and an output host parameter is determined as follows.

Case:

- a) If *PD* is a <status parameter>, then *P* is an output host parameter.
- b) Otherwise,

Case:

- i) If *PN* is contained in an <SQL argument> *A_i* of the <SQL argument list> of a <routine invocation> immediately contained in a <call statement> that is contained in *SPS*, then:
 - 1) Let *R* be the subject routine of the <routine invocation>.
 - 2) Let *PR_i* be the *i*-th SQL parameter of *R*.
 - 3) Case:
 - A) If *PN* is contained in a <host parameter specification> that is the <target specification> that is simply contained in *A_i* and *PR_i* is an output SQL parameter, then *P* is an output host parameter.
 - B) If *PN* is contained in a <host parameter specification> that is the <target specification> that is simply contained in *A_i* and *PR_i* is both an input SQL parameter and an output SQL parameter, then *P* is both an input host parameter and an output host parameter.
 - C) Otherwise, *P* is an input host parameter.
- ii) If *PN* is contained in a <value specification> or a <simple value specification> that is contained in *SPS*, and *PN* is not contained in a <target specification> or a <simple target specification> that is contained in *SPS*, then *P* is an input host parameter.

9.6 Host parameter mode determination

- iii) If PN is contained in a <target specification> or a <simple target specification> that is contained in SPS , and PN is not contained in a <value specification> or a <simple value specification> that is contained in SPS , then P is an output host parameter.
- iv) If PN is contained in a <value specification> or a <simple value specification> that is contained in SPS , and in a <target specification> or a <simple target specification> that is contained in SPS , then P is both an input host parameter and an output host parameter.
- v) Otherwise, P is neither an input host parameter nor an output host parameter.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.7 Type name determination

This Subclause is modified by [Subclause 10.5](#), “Type name determination”, in ISO/IEC 9075-14.

Subclause Signature

```
"Type name determination" [Syntax Rules] (
Parameter: "PREDEF TYPE"
) Returns: "IDENTIFIER"
```

Function

Determine an <identifier> given the name of a predefined data type.

Syntax Rules

- 1) Let *DT* be the *PREDEF TYPE* in an application of this Subclause. The *FNSDT* that is the result of the application of this Subclause is *IDENTIFIER*.
- 2) 14 *FNSDT* is defined as follows.

Case:

- a) If *DT* specifies CHARACTER, then let *FNSDT* be “CHAR”.
- b) If *DT* specifies CHARACTER VARYING, then let *FNSDT* be “VARCHAR”.
- c) If *DT* specifies CHARACTER LARGE OBJECT, then let *FNSDT* be “CLOB”.
- d) If *DT* specifies BINARY, then let *FNSDT* be “BINARY”.
- e) If *DT* specifies BINARY VARYING, then let *FNSDT* be “VARBINARY”.
- f) If *DT* specifies BINARY LARGE OBJECT, then let *FNSDT* be “BLOB”.
- g) If *DT* specifies SMALLINT, then let *FNSDT* be “SMALLINT”.
- h) If *DT* specifies INTEGER, then let *FNSDT* be “INTEGER”.
- i) If *DT* specifies BIGINT, then let *FNSDT* be “BIGINT”.
- j) If *DT* specifies DECIMAL, then let *FNSDT* be “DECIMAL”.
- k) If *DT* specifies NUMERIC, then let *FNSDT* be “NUMERIC”.
- l) If *DT* specifies REAL, then let *FNSDT* be “REAL”.
- m) If *DT* specifies FLOAT, then let *FNSDT* be “FLOAT”.
- n) If *DT* specifies DOUBLE PRECISION, then let *FNSDT* be “DOUBLE”.
- o) If *DT* specifies DATE, then let *FNSDT* be “DATE”.
- p) If *DT* specifies TIME, then let *FNSDT* be “TIME”.

- q) If *DT* specifies **TIMESTAMP**, then let *FNSDT* be “**TIMESTAMP**”.
- r) If *DT* specifies **INTERVAL**, then let *FNSDT* be “**INTERVAL**”.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.8 Determination of identical values

This Subclause is modified by Subclause 9.5, “Determination of identical values”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.6, “Determination of identical values”, in ISO/IEC 9075-14.

Subclause Signature

```
"Determination of identical values" [General Rules] (  
Parameter: "FIRST VALUE",  
Parameter: "SECOND VALUE"  
)
```

Function

Determine whether two instances of values are identical, that is to say, are occurrences of the same value.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *V1* be the *FIRST VALUE* and let *V2* be the *SECOND VALUE* in an application of this Subclause.

NOTE 238 — This Subclause is invoked implicitly wherever the word *identical* is used of two values.

- 2) Case:

- a) If *V1* and *V2* are both the null value, then *V1* is identical to *V2*.
- b) If *V1* is the null value and *V2* is not the null value, or if *V1* is not the null value and *V2* is the null value, then *V1* is not identical to *V2*.
- c) If *V1* and *V2* are of comparable predefined types, then

Case:

- i) If *V1* and *V2* are character strings, then let *L* be CHARACTER_LENGTH(*V1*).

Case:

- 1) If CHARACTER_LENGTH(*V2*) equals *L*, and if for all *i*, $1 \text{ (one)} \leq i \leq L$, the *i*-th character of *V1* corresponds to the same character position of [ISO10646] as the *i*-th character of *V2*, then *V1* is identical to *V2*.
- 2) Otherwise, *V1* is not identical to *V2*.

- ii) If $V1$ and $V2$ are TIME WITH TIME ZONE or TIMESTAMP WITH TIME ZONE and are not distinct, and their time zone displacement fields are not distinct, then $V1$ is identical to $V2$.
- iii) Otherwise, $V1$ is identical to $V2$ if and only if $V1$ is not distinct from $V2$.
- d) 09 14 If $V1$ and $V2$ are of constructed types, then
Case:
 - i) If $V1$ and $V2$ are rows and their respective fields are identical, then $V1$ is identical to $V2$.
 - ii) If $V1$ and $V2$ are arrays and have the same cardinality and elements in the same ordinal position in the two arrays are identical, then $V1$ is identical to $V2$.
 - iii) If $V1$ and $V2$ are multisets and have the same cardinality N and there exist enumerations $VE1_i$, $1 \text{ (one)} \leq i \leq N$, of $V1$ and $VE2_i$, $1 \text{ (one)} \leq i \leq N$, of $V2$ such that for all i , $VE1_i$ is identical to $VE2_i$, then $V1$ is identical to $V2$.
 - iv) If $V1$ and $V2$ are references and $V1$ is not distinct from $V2$, then $V1$ is identical to $V2$.
 - v) Otherwise, $V1$ is not identical to $V2$.
- e) If $V1$ and $V2$ are of the same most specific type MST and MST is a user-defined type, then
Case:
 - i) If MST is a distinct type whose source type is SDT and the results of $SDT(V1)$ and $SDT(V2)$ are identical, then $V1$ is identical to $V2$.
 - ii) If MST is a structured type and, for every observer function O defined for MST , the results of the invocations $O(V1)$ and $O(V2)$ are identical, then $V1$ is identical to $V2$.
 - iii) Otherwise, $V1$ is not identical to $V2$.
- f) Otherwise, $V1$ is not identical to $V2$.

Conformance Rules

None.

9.9 Equality operations

This Subclause is modified by Subclause 9.6, “Equality operations”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.8, “Equality operations”, in ISO/IEC 9075-14.

Function

Specify the prohibitions and restrictions by data type on operations that involve testing for equality.

Syntax Rules

- 1) An *equality operation* is any of the following:
 - a) A <comparison predicate> that specifies <equals operator> or <not equals operator>.
 - b) A <quantified comparison predicate> that specifies <equals operator> or <not equals operator>.
 - c) An <in predicate>.
 - d) A <like predicate>.
 - e) A <similar predicate>.
 - f) A <distinct predicate>.
 - g) A <match predicate>.
 - h) A <member predicate>.
 - i) A <joined table> that specifies NATURAL or USING.
 - j) A <user-defined ordering definition> that specifies MAP.
 - k) A <position expression>.
- 2) An *operand of an equality operation* is any of the following:
 - a) A field of the declared row type of a <row value predicand> that is simply contained in a <comparison predicate> that specifies <equals operator> or <not equals operator>.
 - b) A field of the declared row type of a <row value predicand> that is simply contained in a <quantified comparison predicate> that specifies <equals operator> or <not equals operator>.
 - c) A column of a <table subquery> that is simply contained in a <quantified comparison predicate> that specifies <equals operator> or <not equals operator>.
 - d) A field of the declared row type of a <row value predicand> or <row value expression> that is simply contained in an <in predicate>.
 - e) A column of a <table subquery> that is simply contained in an <in predicate>.
 - f) A field of the declared row type of a <row value predicand> that is simply contained in a <like predicate>.
 - g) A <character pattern>, <escape character>, <octet pattern> or <escape octet> that is simply contained in a <like predicate>.

- h) A field of the declared row type of a <row value predicand> that is simply contained in a <similar predicate>.
 - i) A <similar pattern> or <escape character> that is simply contained in a <similar predicate>.
 - j) A field of the declared row type of a <row value predicand> that is simply contained in a <distinct predicate>.
 - k) A field of the declared row type of a <row value predicand> that is simply contained in a <match predicate>.
 - l) A column of a <table subquery> that is simply contained in a <match predicate>.
 - m) A field of the declared row type of a <row value predicand> that is simply contained in a <member predicate>.
 - n) A corresponding join column of a <joined table> that specifies NATURAL or USING.
 - o) The <returns data type> of the SQL-invoked function identified by a <map function specification> simply contained in a <user-defined ordering definition> that specifies MAP.
 - p) A <string value expression> that is simply contained in a <position expression>.
 - q) A <binary value expression> that is simply contained in a <position expression>.
- 3) The declared type of an operand of an equality operation shall not be UDT-NC-ordered.
 - 4) If the declared type of the operands of an equality operation is a character string type, then the Syntax Rules of [Subclause 9.13, “Collation determination”](#), apply.
 - 5) If the declared type of an operand *OP* of an equality operation is a multiset type, then *OP* is a multiset operand of a multiset element grouping operation. The Syntax Rules of [Subclause 9.11, “Multiset element grouping operations”](#), apply.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared type of an operand of an equality operation shall not be ST-ordered.
- 2) Without Feature T042, “Extended LOB data type support”, in conforming SQL language, the declared type of an operand of an equality operation shall not be LOB-ordered.
- 3) Without Feature S275, “Advanced multiset support”, in conforming SQL language, the declared type of an operand of an equality operation shall not be multiset-ordered.

9.9 Equality operations

NOTE 239 — If the declared type of an operand *OP* of an equality operation is a multiset type, then *OP* is a multiset operand of a multiset element grouping operation. The Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, apply.

9.10 Grouping operations

*This Subclause is modified by Subclause 9.7, “Grouping operations”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 10.9, “Grouping operations”, in ISO/IEC 9075-14.*

Function

Specify the prohibitions and restrictions by data type on operations that involve grouping of data.

Syntax Rules

- 1) A *grouping operation* is any of the following:
 - a) A <group by clause>.
 - b) A <window partition clause>.
 - c) An <aggregate function> that specifies DISTINCT.
 - d) A <query specification> that immediately contains DISTINCT.
 - e) A <query expression body> that simply contains or implies UNION DISTINCT.
 - f) A <query expression body> that simply contains EXCEPT.
 - g) A <query term> that simply contains INTERSECT.
 - h) A <unique predicate>.
 - i) A <unique constraint definition>.
 - j) A <referential constraint definition>.
- 2) An *operand of a grouping operation* is any of the following:
 - a) A grouping column of a <group by clause>.
 - b) A partitioning column of a <window partition clause>.
 - c) A <value expression> simply contained in an <aggregate function> that specifies DISTINCT.
 - d) A column of the result of a <query specification> that immediately contains DISTINCT.
 - e) A column of the result of a <query expression body> that simply contains or implies UNION DISTINCT.
 - f) A column of the result of a <query expression body> that simply contains EXCEPT.
 - g) A column of the result of a <query term> that simply contains INTERSECT.
 - h) A column of the <table subquery> simply contained in a <unique predicate>.
 - i) A column identified by the <unique column list> of a <unique constraint definition>.
 - j) A referencing column of a <referential constraint definition>.
- 3) The declared type of an operand of a grouping operation shall not be LOB-ordered, array-ordered, multiset-ordered, UDT-EC-ordered, or UDT-NC-ordered.

9.10 Grouping operations

- 4) If the declared type of an operand of a grouping operation is a character string type, then the Syntax Rules of [Subclause 9.13](#), “Collation determination”, apply.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared type of an operand of a grouping operation shall not be ST-ordered.

9.11 Multiset element grouping operations

This Subclause is modified by Subclause 9.8, “Multiset element grouping operations”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.10, “Multiset element grouping operations”, in ISO/IEC 9075-14.

Function

Specify the prohibitions and restrictions by data type on the declared element type of a multiset for operations that involve grouping the elements of a multiset.

Syntax Rules

- 1) A *multiset element grouping operation* is any of the following:
 - a) An equality operation such that the declared type of an operand of the equality operation is a multiset type.
 - b) A <multiset set function>.
 - c) A <multiset value expression> that specifies MULTiset UNION DISTINCT.
 - d) A <multiset value expression> that specifies MULTiset EXCEPT.
 - e) A <multiset term> that specifies MULTiset INTERSECT.
 - f) A <submultiset predicate>.
 - g) A <set predicate>.
 - h) A <general set function> that specifies INTERSECTION.
- 2) A *multiset operand* of a multiset element grouping operation is any of the following:
 - a) A <multiset value expression> simply contained in a <multiset set function>.
 - b) A <multiset value expression> or a <multiset term> that is simply contained in a <multiset value expression> that simply contains MULTiset UNION DISTINCT.
 - c) A <multiset value expression> or a <multiset term> that is simply contained in a <multiset value expression> that simply contains MULTiset EXCEPT.
 - d) A <multiset term> or a <multiset primary> that is simply contained in a <multiset term> that simply contains MULTiset INTERSECT.
 - e) A field of a comparand of a <comparison predicate> such that the <comparison predicate> specifies <equals operator> or <not equals operator> and such that the declared type of the field is a multiset type.
 - f) A <multiset value expression> simply contained in a <member predicate>.
 - g) A field of the <row value expression> simply contained in a <submultiset predicate>.
 - h) The <multiset value expression> simply contained in a <submultiset predicate>.
 - i) A field of the <row value expression> simply contained in a <set predicate>.

9.11 Multiset element grouping operations

- j) A value expression> simply contained in a <general set function> that specifies INTERSECTION.
- 3) The declared element type of a multiset operand of a multiset element grouping operation shall not be LOB-ordered, array-ordered, multiset-ordered, UDT-EC-ordered, or UDT-NC-ordered.
- 4) If the declared element type of a multiset operand of a multiset element grouping operation is a character string type, then the Syntax Rules of [Subclause 9.13](#), “[Collation determination](#)”, apply.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared element type of a multiset operand of a multiset element grouping operation shall not be ST-ordered.

9.12 Ordering operations

This Subclause is modified by Subclause 9.9, “Ordering operations”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 10.11, “Ordering operations”, in ISO/IEC 9075-14.

Function

Specify the prohibitions and restrictions by data type on operations that involve ordering of data.

Syntax Rules

- 1) An *ordering operation* is any of the following:
 - a) A <comparison predicate> that does not specify <equals operator> or <not equals operator>.
 - b) A <quantified comparison predicate> that does not specify <equals operator> or <not equals operator>.
 - c) A <between predicate>.
 - d) An <overlaps predicate>.
 - e) An <aggregate function> that specifies MAX or MIN.
 - f) A <sort specification list>.
 - g) A <user-defined ordering definition> that specifies ORDER FULL BY MAP.
- 2) An *operand of an ordering operation* is any of the following:
 - a) A field of the declared row type of a <row value predicand> that is simply contained in a <comparison predicate> that does not specify <equals operator> or <not equals operator>.
 - b) A field of the declared row type of a <row value predicand> that is simply contained in a <quantified comparison predicate> that does not specify <equals operator> or <not equals operator>.
 - c) A column of the <table subquery> that is simply contained in a <quantified comparison predicate> that does not specify <equals operator> or <not equals operator>.
 - d) A field of the declared row type of a <row value predicand> that is simply contained in a <between predicate>.
 - e) A field of the declared row type of a <row value predicand> that is simply contained in an <overlaps predicate>.
 - f) A <value expression> simply contained in an <aggregate function> that specifies MAX or MIN.
 - g) A <value expression> simply contained in a <sort key>.
 - h) The <returns data type> of the SQL-invoked function identified by a <map function specification> simply contained in a <user-defined ordering definition> that specifies ORDER FULL BY MAP.
- 3) The declared type of an operand of an ordering operation shall not be LOB-ordered, array-ordered, multiset-ordered, reference-ordered, UDT-EC-ordered, or UDT-NC-ordered.

- 4) If the declared type of an operand of an ordering operation is a character string type, then the Syntax Rules of [Subclause 9.13](#), “[Collation determination](#)”, apply.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared type of an operand of an ordering operation shall not be ST-ordered.

9.13 Collation determination

Subclause Signature

```
"Collation determination" [Syntax Rules] (  
Parameter: "VALSET"  
) Returns: "COLL"
```

Function

Specify rules for determining the collation to be used in the comparison of character strings.

Syntax Rules

- 1) Let *VS* be the *VALSET* in an application of this Subclause. The collation to be used that is the result of the application of this Subclause is *COLL*.
- 2) Let *TYPESET* be the set of the declared types of all values in *VS*. At least one data type in *TYPESET* shall have a declared type collation.
- 3) The Syntax Rules of [Subclause 9.3, "Result of data type combinations"](#), are applied with *TYPESET* as *DTSET*, yielding the character set *CCS* of the result returned from *RESTYPE*.
- 4) Case:
 - a) If the comparison operation is a <referential constraint definition>, then, for each referencing column, the collation to be used is the declared type collation of the corresponding column of the referenced table.
 - b) If at least one operand has an *explicit* collation derivation, then every operand whose collation derivation is *explicit* shall have the same declared type collation *EDTC* and the collation to be used is *EDTC*.
 - c) If the comparison operation is contained in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement>, or in a <direct SQL statement> that is invoked directly, and *CCS* has an SQL-session collation, then the collation to be used is that SQL-session collation.
 - d) If *CCS* has an SQL-client module collation, then the collation to be used is that collation.
 - e) Otherwise, every operand whose collation derivation is *implicit* shall have the same declared type collation *IDTC* and the collation to be used is *IDTC*.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.14 Execution of array-returning functions

This Subclause is modified by Subclause 8.5, “Execution of array-returning functions”, in ISO/IEC 9075-13.

Subclause Signature

```
"Execution of array-returning functions" [General Rules] (  
Parameter: "ARRAY",  
Parameter: "EFFECTIVE SQL PARAMETER LIST",  
Parameter: "PROGRAM"  
)
```

Function

Define the execution of an external function that returns an array value.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *AR* be the *ARRAY*, let *ESPL* be the *EFFECTIVE SQL PARAMETER LIST*, and let *P* be the *PROGRAM* in an application of this Subclause.
- 2) Let *ARC* be the cardinality of *AR*.
- 3) Let *EN* be the number of entries in *ESPL*.
- 4) Let *ESP_i*, 1 (one) $\leq i \leq EN$, be the *i*-th parameter in *ESPL*.
- 5) Let *FRN* be the number of result data items.
- 6) 13 Let *PN* and *N* be the number of values in the static SQL argument list of *P*.
- 7) Let *E* be 0 (zero).
- 8) 13 If the call type data item has a value of -1 (indicating “open call”), then *P* is executed with a list of *EN* parameters *PD_i* whose parameter names are *PN_i* and whose values are set as follows:
 - a) Depending on whether the language of *R* specifies ADA, C, COBOL, FORTRAN, M, PASCAL, or PLI, let the *operative data type correspondences* table be Table 16, “Data type correspondences for Ada”, Table 17, “Data type correspondences for C”, Table 18, “Data type correspondences for COBOL”, Table 19, “Data type correspondences for Fortran”, Table 20, “Data type correspondences for M”, Table 21, “Data type correspondences for Pascal”, or Table 22, “Data type correspondences for PL/I”,

9.14 Execution of array-returning functions

respectively. Refer to the two columns of the operative data type correspondences table as the “SQL data type” column and the “host data type” column.

- b) For i varying from 1 (one) to EN , the data type DT_i of PD_i is the data type listed in the host data type column of the row in the data type correspondences table whose value in the SQL data type column corresponds to the data type of ESP_i .
- c) The value of PD_i is set to the value of ESP_i .

9) 13 Case:

- a) 13 If the value of the exception data item is '00000' (corresponding to the completion condition *successful completion*) or the first 2 characters are '01' (corresponding to the completion condition *warning* with any subcondition), then set the call type data item to 0 (zero) (indicating *fetch call*).
- b) 13 If the exception data item is '02000' (corresponding to the completion condition *no data*):
 - i) If each PD_i , for i ranging from $(PN+FRN)+N+1$ through $(PN+FRN)+N+FRN$ (that is, the SQL indicator arguments corresponding to the result data items), is negative, then set AR to the null value.
 - ii) Set the call type data item to 1 (one) (indicating *close call*).
- c) 13 Otherwise, set the call type data item to 1 (one) (indicating *close call*).

10) The following steps are applied as long as the call type data item has a value 0 (zero) (corresponding to *fetch call*):

- a) 13 P is executed with a list of EN parameters PD_i whose parameter names are PN_i and whose values are set as follows:
 - i) For i varying from 1 (one) to EN , the <data type> DT_i of PD_i is the data type listed in the host data type column of the row in the data type correspondences table whose value in the SQL data type column corresponds to the data type of ESP_i .
 - ii) For i ranging from 1 (one) to $EN-2$, the value of PD_i is set to the value of ESP_i .
 - iii) For the save area data item, for i equal to $EN-1$, the value of PD_i is set to the value returned in PD_i by the prior execution of P .
 - iv) For the call type data item, for i equal to EN , the value of PD_i is set to 0 (zero).
- b) 13 Case:
 - i) 13 If the exception data item is '00000' (corresponding to completion condition *successful completion*) or the first 2 characters are '01' (corresponding to completion condition *warning* with any subcondition), then:
 - 1) Increment E by 1 (one).
 - 2) If $E > ARC$, then an exception condition is raised: *data exception — array element error*.
 - 3) If the call type data item is 0 (zero), then

Case:

9.14 Execution of array-returning functions

A) 13 If each PD_i , for i ranging from $(PN+FRN)+N+1$ through $(PN+FRN)+N+FRN$ (that is, the SQL indicator arguments corresponding to the result data items) is negative, then let the E -th element of AR be the null value.

B) Otherwise,

Case:

I) 13 If FRN is 1 (one), then let the E -th element of AR be the value of the result data item.

II) Otherwise:

1) Let RDI_i , $1 \text{ (one)} \leq i \leq FRN$, be the value of the i -th result data item.

2) Let the E -th element of AR be the value of the following <row value expression>:

ROW (RDI_1 , . . . ,
 RDI_{FRN})

ii) 13 If the exception data item is '02000' (corresponding to completion condition *no data*), then:

1) If the value of E is 0 (zero), then set AR to an array whose cardinality is 0 (zero).

2) Set the call type data item to 1 (one) (indicating *close call*).

iii) Otherwise, set the value of the call type data item to 1 (one) (indicating *close call*).

11) 13 If the call type data item has a value of 1 (one) (indicating *close call*), then P is executed with a list of EN parameters PD_i whose parameter names are PN_i and whose values are set as follows:

a) For i varying from 1 (one) to EN , the <data type> DT_i of PD_i is the data type listed in the host data type column of the row in the data type correspondences table whose value in the SQL data type column corresponds to the data type of ESP_i .

b) For i ranging from 1 (one) to $EN-2$, the value of PD_i is set to the value of ESP_i .

c) For the save area data item, for i equal to $EN-1$, the value of PD_i is set to the value returned in PD_i by the prior execution of P .

d) For the call type data item, for i equal to EN , the value of PD_i is set to 1 (one).

Conformance Rules

None.

9.15 Execution of multiset-returning functions

Subclause Signature

```
"Execution of multiset-returning functions" [General Rules] (  
Parameter: "MULTISET" ,  
Parameter: "EFFECTIVE SQL PARAMETER LIST" ,  
Parameter: "PROGRAM"  
)
```

Function

Define the execution of an external function that returns a multiset value.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *MU* be the *MULTISET*, let *ESPL* be the *EFFECTIVE SQL PARAMETER LIST*, and let *P* be the *PROGRAM* in an application of this Subclause.
- 2) Let *ET* be the element type of *MU*.
- 3) Let *C* be the maximum implementation-defined cardinality of array type with element type *ET*.
- 4) Let *AT* be the array type *ET* ARRAY[*C*].
- 5) Let *AR* be an array whose declared type is *AT*.
- 6) The General Rules of Subclause 9.14, "Execution of array-returning functions", are applied with *AR* as *ARRAY*, *ESPL* as *EFFECTIVE SQL PARAMETER LIST*, and *P* as *PROGRAM*.
- 7) Let *MU* be the result of casting *AR* to the multiset type of *MU* according to the General Rules of Subclause 6.12, "<cast specification>".

Conformance Rules

None.

9.16 XQuery regular expression matching

Subclause Signature

```
"XQuery regular expression matching" [General Rules] (  
Parameter: "STRING" ,  
Parameter: "PATTERN" ,  
Parameter: "POSITION" ,  
Parameter: "UNITS" ,  
Parameter: "FLAG"  
) Returns: "LIST"
```

Function

Determine a list of match vectors to an XQuery regular expression in a character string.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *S* be the *STRING*, let *PAT* be the *PATTERN*, let *SP* be the *POSITION*, let *U* be the *UNITS*, and let *FL* be the *FLAG* in an application of this Subclause. The list of match vectors that is the result of the application of this Subclause is *LIST*.
- 2) Case:
 - a) If the character repertoire of *PAT* is UCS, then let *P* be *PAT*.
 - b) Otherwise, let *P* be the result of an implementation-defined conversion of *PAT* to UCS.
- 3) Case:
 - a) If the character repertoire of *FL* is UCS, then let *F* be *FL*.
 - b) Otherwise, let *F* be the result of an implementation-defined conversion of *FL* to UCS.
- 4) If *P* is not an XQuery regular expression, then an exception condition is raised: *data exception — invalid XQuery regular expression*.
- 5) If *F* is not an XQuery option flag, then an exception condition is raised: *data exception — invalid XQuery option flag*.
- 6) Let *LEN* be the length in characters of *S*.

9.16 XQuery regular expression matching

- 7) A position/length PL is a pair of exact numeric values, a position P and a length L , such that either $P = 0$ (zero) and $L = 0$ (zero), or P is positive, L is non-negative, and $P + L \leq LEN + 1$ (one). If P is positive, then PL represents the substring of S whose starting position is P and whose length is L . If P is 0 (zero), then PL does not represent a substring of S .

NOTE 240 — For example, if S is 'a', then there are four position/lengths, namely (1,0), denoting the zero-length substring at the beginning of S ; (2,0), denoting the zero-length substring at the end of S ; (1,1), denoting the whole string S ; and (0,0), denoting no substring of S .

- 8) Let G be the number of parenthesized groups in PAT .
- 9) An XQuery regular expression match is defined as follows:

Case:

- a) If the character repertoire of S is UCS, then an XQuery regular expression match is a match as defined by [XQueryFO] for the function `fn:matches()`, with F as the value of the parameter `$flags`, with the following modifications:
- i) The metacharacters “^”, “\$”, and “.” are interpreted according to the guidelines of RL1.6 “Line boundaries”, in [Unicode18].
 - ii) The multi-character escape “\s” matches U+0020 (space), U+0009 (tab), or any single character or two-character sequence identified as a newline sequence by RL1.6 “line boundaries” in [Unicode18].
- b) Otherwise, the definition of XQuery regular expression match is implementation-defined.

- 10) A match vector MV is an array of $G + 1$ position/lengths, indexed from 0 (zero) through G . MV represents an XQuery regular expression match of PAT , including information about capture groups. $MV[0]$ is the position/length that denotes the entire XQuery regular expression match of PAT . For i between 1 (one) and G , $MV[i]$ is the position/length that denotes the i -th capture group within the XQuery regular expression match that is denoted by $MV[0]$. If the i -th parenthesized subexpression of PAT is not matched, then the position of $MV[i]$ is 0 (zero).

NOTE 241 — The position of $MV[0]$ is never 0 (zero), since an XQuery regular expression match to PAT can never be “no substring”. An example of a capture group that is not matched occurs when S is 'a' and PAT is '(a)|(b)'. In this example, the match vector is given by $MV[0] = (1,1)$, $MV[1] = (1,1)$, and $MV[2] = (0,0)$, since the second parenthesized subexpression has no match.

- 11) The result of this Subclause is a list $LIST$ of match vectors, defined as follows:

- a) Let P_0 be $SP - 1$ (one), and let L_0 be 0 (zero).
- b) For $i \geq 1$ (one), P_i and L_i are defined recursively as follows: Let P_i be the position in units U of the first XQuery regular expression match of PAT in S such that $P_i > P_{i-1} + L_{i-1}$, and let L_i be the length in units U of this match.

NOTE 242 — The sequence P_i is strictly increasing and bounded above by $LEN + 1$ (one), so this recursive definition terminates after finding a finite number of disjoint matches.

- c) Let N be the number of XQuery regular expression matches found.

NOTE 243 — It is possible that there are no matches, in which case N is 0 (zero), and $LIST$ is empty. $(P_0, L_0) = (SP - 1, 0)$ is not included in the result; it is merely used as the position to start seeking the first match.

- d) For i between 1 (one) and N , let MV_i be the match vector such that:
- i) $MV_i[0]$ is the position/length whose position is P_i and whose length is L_i .

ii) For j between 1 (one) and G , $MV_i[j]$ is the position/length

Case:

- 1) If the j -th parenthesized sub-expression of PAT is not matched in the i -th XQuery regular expression match, then position 0 (zero) and length 0 (zero).
- 2) Otherwise, the position and length of the j -th capture group within the i -th XQuery regular expression match.

e) Let $LIST$ be the list of N match vectors MV_i in order from 1 (one) to N .

12) $LIST$ is the result of this Subclause.

Conformance Rules

None.

9.17 XQuery regular expression replacement

Subclause Signature

```
"XQuery regular expression replacement" [General Rules] (
Parameter: "STRING" ,
Parameter: "PATTERN" ,
Parameter: "MATCH" ,
Parameter: "REPLACEMENT" ,
Parameter: "FLAG"
) Returns: "RESULT"
```

Function

Replace a substring that matches an XQuery regular expression with a replacement character string.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *STR* be the *STRING*, let *PAT* be the *PATTERN*, let *MV* be the *MATCH*, let *REP* be the *REPLACEMENT*, and let *FL* be the *FLAG* in an application of this Subclause. The character string that is the result of the application of this Subclause is *RESULT*.
- 2) Let *PL* be *MV*[0]. Let *P* be the position of *PL* and let *L* be the length of *PL*.
- 3) If *L* is zero, then an exception condition is raised: *data exception — attempt to replace a zero-length string*.
- 4) Case:
 - a) If the character repertoire of *REP* is UCS, then let *R* be *REP*.
 - b) Otherwise, let *R* be the result of an implementation-defined conversion of *REP* to UCS.
- 5) If *R* is not an XQuery replacement string, then an exception condition is raised: *data exception — invalid XQuery replacement string*.
- 6) Let *STR*₁ be the character string consisting of the first (*P*−1) characters of *STR*.
- 7) Let *STR*₂ be the character string consisting of the *P*-th through (*P*+*L*−1)-th character of *STR*.
- 8) Let *STR*₃ be the character string consisting of the (*P*+*L*)-th character through the last character of *STR*.
- 9) Case:

9.17 XQuery regular expression replacement

- a) If the character repertoire of *STR*, *PAT*, *REP*, and *FL* is UCS, then let *T* be the result of the XQuery function `fn:replace()`, where the value of the `$input` argument is *STR*₂, the value of the `$pattern` argument is *PAT*, the value of the `$replacement` argument is *R*, and the value of the `$flags` argument is *FL*.
 - b) Otherwise, the value of *T* is implementation-defined.
- 10) The result of this Subclause is the concatenation of *STR*₁, *T*, and *STR*₃.

Conformance Rules

None.

9.18 Data type identity

Function

Determine whether two data types are compatible and have the same characteristics.

Syntax Rules

- 1) Let PM and P be the two data types specified in an application of this Subclause.
- 2) PM and P shall be compatible.
- 3) If PM is a fixed-length character string type, then the length of PM shall be equal to the length of P .
- 4) If PM is a variable-length character string type, then the maximum length of PM shall be equal to the maximum length of P .
- 5) If PM is an exact numeric type, then the precision and scale of PM shall be equal to the precision and scale of P , respectively.
- 6) If PM is an approximate numeric type, then the precision of PM shall be equal to the precision of P .
- 7) If PM is a fixed-length binary string type, then the length of PM shall be equal to the length of P .
- 8) If PM is a variable-length binary string type, then the maximum length of PM shall be equal to the maximum length of P .
- 9) If PM is a datetime data type with <time fractional seconds precision>, then the <time fractional seconds precision> of PM shall be equal to the <time fractional seconds precision> of P .
- 10) If PM is an interval type, then the <interval qualifier> of PM shall be equivalent to the <interval qualifier> of P .
- 11) If PM is a collection type, then:
 - a) The kind of collection (ARRAY or MULTISSET) of PM and the kind of collection of P shall be the same.
 - b) If PM is an array type, then the maximum cardinality of PM shall be equal to the the maximum cardinality of P .
 - c) The Syntax Rules of this Subclause are applied with the element type of PM and the element type of P as the two data types.
- 12) If PM is a row type, then:
 - a) Let N be the degree of PM .
 - b) Let $DTFPM_i$ and $DTFP_i$, $1 \text{ (one)} \leq i \leq N$, be the data type of the i -th field of PM and of P , respectively. For i varying from 1 (one) to N , the Syntax Rules of this Subclause are applied with $DTFPM_i$ and $DTFP_i$ the two data types.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.19 Determination of a from-sql function

Function

Determine the from-sql function of a user-defined type given the name of a user-defined type and the name of the group.

Syntax Rules

- 1) Let *UDT* and *GN* be a *TYPE* and a *GROUP* specified in an application of this Subclause.
- 2) Let *SSUDT* be the set of supertypes of *UDT*.
- 3) Let *SUDT* be the data type, if any, in *SSUDT* such that the transform descriptor included in the data type descriptor of *SUDT* includes a group descriptor *GD* that includes a group name that is equivalent to *GN*.
- 4) The *applicable from-sql function* is the SQL-invoked function identified by the specific name of the from-sql function, if any, in *GD*.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.20 Determination of a from-sql function for an overriding method

Function

Determine the from-sql function of a user-defined type given the name of an overriding method and the ordinal position of an SQL parameter.

Syntax Rules

- 1) Let *R* and *N* be a *ROUTINE* and a *POSITION* specified in an application of this Subclause.
- 2) Let *OM* be original method of *R*.
- 3) The *applicable from-sql function* is the from-sql function associated with the *N*-th SQL parameter of *OM*, if any.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.21 Determination of a to-sql function

Function

Determine the to-sql function of a user-defined type given the name of a user-defined type and the name of a group.

Syntax Rules

- 1) Let *UDT* and *GN* be a *TYPE* and a *GROUP* specified in an application of this Subclause.
- 2) Let *SSUDT* be the set of supertypes of *UDT*.
- 3) Let *SUDT* be the data type, if any, in *SSUDT* such that the transform descriptor included in the data type descriptor of *SUDT* includes a group descriptor *GD* that includes a group name that is equivalent to *GN*.
- 4) The *applicable to-sql function* is the SQL-invoked function identified by the specific name of the to-sql function, if any, in *GD*.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.22 Determination of a to-sql function for an overriding method

Function

Determine the to-sql function of a user-defined type given the name of an overriding method.

Syntax Rules

- 1) Let *R* be a *ROUTINE* specified in an application of this Subclause.
- 2) Let *OM* be the original method of *R*
- 3) The *applicable to-sql function* is the SQL-invoked function associated with the result of *OM*, if any.

Access Rules

None.

General Rules

None.

Conformance Rules

None.

9.23 Generation of the next value of a sequence generator

Function

Generate and return the next value of a sequence generator.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *SEQ* be the *SEQUENCE* specified in an application of this Subclause.
- 2) Let *DT*, *CBV*, *INC*, *SMAX*, and *SMIN* be the data type, current base value, increment, maximum value and minimum value, respectively, of *SEQ*.
- 3) If there exists a non-negative integer *N* such that $SMIN \leq CBV + N * INC \leq SMAX$ and the value $(CBV + N * INC)$ has not already been returned in the current cycle, then let *VI* be $(CBV + N * INC)$. Otherwise,
Case:
 - a) If the cycle option of *SEQ* is NO CYCLE, then an exception condition is raised: *data exception — sequence generator limit exceeded*.
 - b) Otherwise, a new cycle is initiated.
Case:
 - i) If *SEQ* is an ascending sequence generator, then let *VI* be *SMIN*.
 - ii) Otherwise, let *VI* be *SMAX*.
- 4) Case:
 - a) If *SEQ* is an ascending sequence generator, the current base value of *SEQ* is set to the value of the lowest non-issued value in the cycle.
 - b) Otherwise, the current base value of *SEQ* is set to the highest non-issued value in the cycle.
- 5) *VI* is returned as the *RESULT*.

Conformance Rules

None.

9.24 Creation of a sequence generator

Function

Complete the definition of an external or internal sequence generator.

Syntax Rules

- 1) Let *OPT* and *DT* be the *OPTIONS* and *DATA TYPE* specified in an application of this Subclause. *OPT* shall conform to the Format of <common sequence generator options>. The BNF nonterminal symbols used in the remainder of this Subclause refer to the contents of *OPT*.
- 2) Each of <sequence generator start with option>, <sequence generator increment by option>, <sequence generator maxvalue option>, <sequence generator minvalue option>, and <sequence generator cycle option> shall be specified at most once.
- 3) If <sequence generator increment by option> is specified, then let *INC* be <sequence generator increment>; otherwise, let *INC* be a <signed numeric literal> whose value is 1 (one).
- 4) The value of *INC* shall not be 0 (zero).
- 5) If the value of *INC* is negative, then *SEQ* is a *descending sequence generator*; otherwise, *SEQ* is an *ascending sequence generator*.
- 6) Case:
 - a) If <sequence generator maxvalue option> is specified, then
Case:
 - i) If NO MAXVALUE is specified, then let *SMAX* be an implementation-defined <signed numeric literal> of declared type *DT*.
 - ii) Otherwise, let *SMAX* be <sequence generator max value>.
 - b) Otherwise, let *SMAX* be an implementation-defined <signed numeric literal> of declared type *DT*.
- 7) Case:
 - a) If <sequence generator minvalue option> is specified, then
Case:
 - i) If NO MINVALUE is specified, then let *SMIN* be an implementation-defined <signed numeric literal> of declared type *DT*.
 - ii) Otherwise, let *SMIN* be <sequence generator min value>.
 - b) Otherwise, let *SMIN* be an implementation-defined <signed numeric literal> of declared type *DT*.
- 8) Case:
 - a) If <sequence generator start with option> is specified, then let *START* be <sequence generator start value>.
 - b) Otherwise,

9.24 Creation of a sequence generator

Case:

- i) If *SEQ* is an ascending sequence generator, then let *START* be *SMIN*.
 - ii) Otherwise, let *START* be *SMAX*.
- 9) The values of *INC*, *START*, *SMAX*, and *SMIN* shall all be exactly representable with the precision and scale of *DT*.
 - 10) The value of *SMAX* shall be greater than the value of *SMIN*.
 - 11) The value of *START* shall be greater than or equal to the value of *SMIN* and lesser than or equal to the value of *SMAX*.
 - 12) If <sequence generator cycle option> is specified, then let *CYC* be <sequence generator cycle option>; otherwise, let *CYC* be NO CYCLE.

Access Rules

None.

General Rules

- 1) A sequence generator descriptor *SEQDS* that describes *SEQ* is created. *SEQDS* includes:
 - a) The sequence generator name that is a zero-length character string.

NOTE 244 — The name of an external sequence generator is later set by GR 1) of Subclause 11.63, “<sequence generator definition>”; however, internal sequence generators are anonymous.
 - b) The data type descriptor of *DT*.
 - c) The increment specified by *INC*.
 - d) The maximum value specified by *SMAX*.
 - e) The minimum value specified by *SMIN*.
 - f) The cycle option specified by *CYC*.
 - g) The current base value, set to *START*.

Conformance Rules

None.

9.25 Altering a sequence generator

Function

Complete the alteration of an internal or external sequence generator.

Syntax Rules

- 1) Let *OPT* and *SEQ* be the *OPTIONS* and *SEQUENCE* specified in an application of this Subclause. *OPT* shall conform to the Format of <alter sequence generator options>. The BNF nonterminal symbols used in the remainder of this Subclause refer to the contents of *OPT*.
- 2) Let *DT* be the data type descriptor included in *SEQ*.
- 3) Each of <alter sequence generator restart option>, <sequence generator increment by option>, <sequence generator maxvalue option>, <sequence generator minvalue option>, and <sequence generator cycle option> shall be specified at most once.
- 4) Case:
 - a) If <sequence generator increment> is specified, then:
 - i) Let *NEWIV* be <sequence generator increment>.
 - ii) The value of *NEWIV* shall not be 0 (zero).
 - b) Otherwise, let *NEWIV* be the increment of *SEQ*.
- 5) Case:
 - a) If <sequence generator maxvalue option> is specified, then
Case:
 - i) If NO MAXVALUE is specified, then let *NEWMAX* be an implementation-defined <signed numeric literal> of declared type *DT*.
 - ii) Otherwise, let *NEWMAX* be <sequence generator max value>.
 - b) Otherwise let *NEWMAX* be the maximum value of *SEQ*.
- 6) Case:
 - a) If <sequence generator minvalue option> is specified, then
Case:
 - i) If NO MINVALUE is specified, then let *NEWMIN* be an implementation-defined <signed numeric literal> of declared type *DT*.
 - ii) Otherwise, let *NEWMIN* be <sequence generator min value>.
 - b) Otherwise let *NEWMIN* be the minimum value of *SEQ*.
- 7) If <sequence generator cycle option> is specified, then let *NEWCYCLE* be <sequence generator cycle option>; otherwise, let *NEWCYCLE* be the cycle option of *SEQ*.

9.25 Altering a sequence generator

- 8) If <alter sequence generator restart option> is specified, then let *NEWVAL* be <sequence generator restart value>; otherwise, let *NEWVAL* be the current base value of *SEQ*.
- 9) The values of *NEWIV*, *NEWMAX*, *NEWMIN*, and *NEWVAL* shall all be exactly representable with the precision and scale of *DT*.
- 10) The value of *NEWMIN* shall be less than the value of *NEWMAX*.
- 11) The value of *NEWVAL* shall be greater than or equal to the value of *NEWMIN* and lesser than or equal to the value of *NEWMAX*.

Access Rules

None.

General Rules

- 1) *SEQ* is modified as follows:
 - a) The increment is set to *NEWIV*.
 - b) The maximum value is set to *NEWMAX*.
 - c) The minimum value is set to *NEWMIN*.
 - d) The cycle option is set to *NEWCYCLE*.
 - e) The current base value is set to *NEWVAL*.

Conformance Rules

None.

10 Additional common elements

10.1 <interval qualifier>

Function

Specify the precision of an interval data type.

Format

```
<interval qualifier> ::=
    <start field> TO <end field>
    | <single datetime field>

<start field> ::=
    <non-second primary datetime field>
    [ <left paren> <interval leading field precision> <right paren> ]

<end field> ::=
    <non-second primary datetime field>
    | SECOND [ <left paren> <interval fractional seconds precision> <right paren> ]

<single datetime field> ::=
    <non-second primary datetime field>
    [ <left paren> <interval leading field precision> <right paren> ]
    | SECOND [ <left paren> <interval leading field precision>
    [ <comma> <interval fractional seconds precision> ] <right paren> ]

<primary datetime field> ::=
    <non-second primary datetime field>
    | SECOND

<non-second primary datetime field> ::=
    YEAR
    | MONTH
    | DAY
    | HOUR
    | MINUTE

<interval fractional seconds precision> ::=
    <unsigned integer>

<interval leading field precision> ::=
    <unsigned integer>
```

Syntax Rules

- 1) There is an ordering of significance of <primary datetime field>s. In order from most significant to least significant, the ordering is: YEAR, MONTH, DAY, HOUR, MINUTE, and SECOND. A <start field> or <single datetime field> with an <interval leading field precision> *i* is more significant than a <start field> or <single datetime field> with an <interval leading field precision> *j* if *i*>*j*. An <end field> or <single datetime field> with an <interval fractional seconds precision> *i* is less significant than an <end field> or <single datetime field> with an <interval fractional seconds precision> *j* if *i*>*j*.
- 2) If TO is specified, then <start field> shall be more significant than <end field> and <start field> shall not specify MONTH. If <start field> specifies YEAR, then <end field> shall specify MONTH.
- 3) The maximum value of <interval leading field precision> is implementation-defined, but shall not be less than 2.
- 4) The maximum value of <interval fractional seconds precision> is implementation-defined, but shall not be less than 6.
- 5) An <interval leading field precision>, if specified, shall be greater than 0 (zero) and shall not be greater than the implementation-defined maximum. If <interval leading field precision> is not specified, then an <interval leading field precision> of 2 is implicit.
- 6) An <interval fractional seconds precision>, if specified, shall be greater than or equal to 0 (zero) and shall not be greater than the implementation-defined maximum. If SECOND is specified and <interval fractional seconds precision> is not specified, then an <interval fractional seconds precision> of 6 is implicit.
- 7) The precision of a field other than the <start field> or <single datetime field> is
Case:
 - a) If the field is not SECOND, then 2.
 - b) Otherwise, 2 digits before the decimal point and the explicit or implicit <interval fractional seconds precision> after the decimal point.

Access Rules

None.

General Rules

- 1) An item qualified by an <interval qualifier> contains the datetime fields identified by the <interval qualifier>.
Case:
 - a) If the <interval qualifier> specifies a <single datetime field>, then the <interval qualifier> identifies a single <primary datetime field>. Any reference to the *most significant* or *least significant* <primary datetime field> of the item refers to that <primary datetime field>.
 - b) Otherwise, the <interval qualifier> identifies those datetime fields from <start field> to <end field>, inclusive.
- 2) An <interval leading field precision> specifies

Case:

- a) If the <primary datetime field> is SECOND, then the number of decimal digits of precision before the specified or implied decimal point of the seconds <primary datetime field>.
 - b) Otherwise, the number of decimal digits of precision of the first <primary datetime field>.
- 3) An <interval fractional seconds precision> specifies the number of decimal digits of precision following the specified or implied decimal point in the <primary datetime field> SECOND.
- 4) The length in positions of an item of type interval is computed as follows.

Case:

- a) If the item is a year-month interval, then

Case:

- i) If the <interval qualifier> is a <single datetime field>, then the length in positions of the item is the implicit or explicit <interval leading field precision> of the <single datetime field>.
 - ii) Otherwise, the length in positions of the item is the implicit or explicit <interval leading field precision> of the <start field> plus 2 (the length of the <non-second primary datetime field> that is the <end field>) plus 1 (one) (the length of the <minus sign> between the <years value> and the <months value> in a <year-month literal>).
- b) Otherwise,

Case:

- i) If the <interval qualifier> is a <single datetime field> that does not specify SECOND, then the length in positions of the item is the implicit or explicit <interval leading field precision> of the <single datetime field>.
- ii) If the <interval qualifier> is a <single datetime field> that specifies SECOND, then the length in positions of the item is the implicit or explicit <interval leading field precision> of the <single datetime field> plus the implicit or explicit <interval fractional seconds precision>. If <interval fractional seconds precision> is greater than zero, then the length in positions of the item is increased by 1 (one) (the length in positions of the <period> between the <seconds integer value> and the <seconds fraction>).
- iii) Otherwise, let *participating datetime fields* mean the datetime fields that are less significant than the <start field> and more significant than the <end field> of the <interval qualifier>. The length in positions of each participating datetime field is 2.

Case:

- 1) If <end field> is SECOND, then the length in positions of the item is the implicit or explicit <interval leading field precision>, plus 3 times the number of participating datetime fields (each participating datetime field has length 2 positions, plus the <minus sign>s or <colon>s that precede them have length 1 (one) position), plus the implicit or explicit <interval fractional seconds precision>, plus 3 (the length in positions of the <end field> other than any <interval fractional seconds precision> plus the length in positions of its preceding <colon>). If <interval fractional seconds precision> is greater than zero, then the length in positions of the item is increased by 1 (one) (the length in positions of the <period> within the field identified by the <end field>).

- 2) Otherwise, the length in positions of the item is the implicit or explicit <interval leading field precision>, plus 3 times the number of participating datetime fields (each participating datetime field has length 2 positions, plus the <minus sign>s or <colon>s that precede them have length 1 (one) position), plus 2 (the length in positions of the <end field>), plus 1 (one) (the length in positions of the <colon> preceding the <end field>).

Conformance Rules

- 1) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval qualifier>.

10.2 <language clause>

This Subclause is modified by Subclause 8.4, “<language clause>”, in ISO/IEC 9075-13.

Function

Specify a programming language.

Format

```
<language clause> ::=
    LANGUAGE <language name>
```

```
13 <language name> ::=
    ADA
    | C
    | COBOL
    | FORTRAN
    | M | MUMPS
    | PASCAL
    | PLI
    | SQL
```

Syntax Rules

- 1) If MUMPS is specified, then M is implicit.

Access Rules

None.

General Rules

- 1) The host language specified by the <language clause> is defined in the International Standard identified by the <language name> keyword. Table 15, “Standard programming languages”, specifies the relationship.

13 Table 15 — Standard programming languages

Language keyword	Relevant standard
ADA	[ISO8652]
C	[ISO9899]
COBOL	[ISO1989]

Language keyword	Relevant standard
FORTRAN	[ISO1539]
M	[ISO11756]
PASCAL	[ISO7185] and [ISO10206]
PLI	[ISO6160]
SQL	ISO/IEC 9075

Conformance Rules

None.

10.3 <path specification>

Function

Specify an order for searching for an SQL-invoked routine.

Format

```
<path specification> ::=  
    PATH <schema name list>  
  
<schema name list> ::=  
    <schema name> [ { <comma> <schema name> }... ]
```

Syntax Rules

- 1) No two <schema name>s contained in <schema name list> shall be equivalent.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <path specification>.

10.4 <routine invocation>

*This Subclause is modified by Subclause 8.1, “<routine invocation>”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 7.1, “<routine invocation>”, in ISO/IEC 9075-10.
This Subclause is modified by Subclause 8.3, “<routine invocation>”, in ISO/IEC 9075-13.
This Subclause is modified by Subclause 11.1, “<routine invocation>”, in ISO/IEC 9075-14.*

Function

Invoke an SQL-invoked routine.

Format

```
<routine invocation> ::=
    <routine name> <SQL argument list>

<routine name> ::=
    [ <schema name> <period> ] <qualified identifier>

<SQL argument list> ::=
    <left paren> [ <SQL argument> [ { <comma> <SQL argument> }... ] ] <right paren>

<SQL argument> ::=
    <value expression>
    | <generalized expression>
    | <target specification>

<generalized expression> ::=
    <value expression> AS <path-resolved user-defined type name>
```

Syntax Rules

- 1) Let *RI* be the *ROUTINE INVOCATION*, let *TP* be the *SQLPATH* (if any), and let *UDTSM* be the *UDT* (if any) specified in an application of this Subclause.
- 2) Let *RN* be the <routine name> immediately contained in *RI*.
- 3) If *RI* is immediately contained in a <call statement>, then the <SQL argument list> of *RI* shall not contain a <generalized expression> without an intervening <routine invocation>.
- 4) Case:
 - a) If *RI* is immediately contained in a <call statement>, then an SQL-invoked routine *R* is a *possibly candidate routine* for *RI* (henceforth, simply “possibly candidate routine”) if *R* is an SQL-invoked procedure and the <qualified identifier> of the <routine name> of *R* is equivalent to the <qualified identifier> of *RN*.
 - b) If *RI* is immediately contained in a <method selection>, then an SQL-invoked routine *R* is a *possibly candidate routine* for *RI* if *R* is an instance SQL-invoked method and the <qualified identifier> of the <routine name> of *R* is equivalent to the <qualified identifier> of *RN*.

- c) If *RI* is immediately contained in a <constructor method selection>, then an SQL-invoked routine *R* is a *possibly candidate routine* for *RI* if *R* is an SQL-invoked constructor method and the <qualified identifier> of the <routine name> of *R* is equivalent to the <qualified identifier> of *RN*.
- d) If *RI* is immediately contained in a <static method selection>, then an SQL-invoked routine *R* is a *possibly candidate routine* for *RI* if *R* is a static SQL-invoked method and the <qualified identifier> of the <routine name> of *R* is equivalent to the <qualified identifier> of *RN* and the method specification descriptor for *R* is included in a user-defined type descriptor for *UDTSM* or for some supertype of *UDTSM*.
- e) Otherwise, an SQL-invoked routine *R* is a *possibly candidate routine* for *RI* if *R* is an SQL-invoked regular function and the <qualified identifier> of the <routine name> of *R* is equivalent to the <qualified identifier> of *RN*.

5) 04 Case:

- a) If *RI* is contained in an <SQL schema statement>, then an SQL-invoked routine *R* is an *executable routine* if and only if *R* is a possibly candidate routine and the applicable privileges for the <authorization identifier> that owns the containing schema include EXECUTE on *R*.
- b) Otherwise, an SQL-invoked routine *R* is an *executable routine* if and only if *R* is a possibly candidate routine and the current privileges include EXECUTE on *R*.

6) Case:

- a) If <SQL argument list> does not immediately contain at least one <SQL argument>, then an *invocable routine* is an executable routine that has no SQL parameters.
- b) Otherwise:
 - i) Let *NA* be the number of <SQL argument>s in the <SQL argument list> *AL* of *RI*. Let *A_i*, 1 (one) ≤ *i* ≤ *NA*, be the *i*-th <SQL argument> in *AL*.
 - ii) Let the *static SQL argument list* of *RI* be *AL*.
 - iii) Let *P_i* be the *i*-th SQL parameter of an executable routine. An *invocable routine* is an SQL-invoked routine *SIR* that is an executable routine such that:
 - 1) *SIR* has *NA* SQL parameters.
 - 2) If *RI* is not immediately contained in a <call statement>, then for each *A_i* that is not a <dynamic parameter specification>,

Case:

A) If the declared type of *P_i* is a user-defined type, then:

- I) Let *ST_i* be the set of subtypes of the declared type of *A_i*.
- II) The Syntax Rules of Subclause 9.5, “Type precedence list determination”, are applied with the data type of some type in *ST_i* as *DATA TYPE*, yielding *TPL_i* returned from *TYPE PRECEDENCE LIST*.
- III) The type designator of the declared type of *P_i* shall be in *TPL_i*.

B) Otherwise:

- I) The Syntax Rules of Subclause 9.5, “Type precedence list determination”, are applied with the declared type of A_i as *DATA TYPE*, yielding TPL_i returned from *TYPE PRECEDENCE LIST*.
 - II) The type designator of the declared type of P_i shall be in TPL_i .
- 7) If <SQL argument list> does not immediately contain at least one <SQL argument>, then:
- a) Let AL be an empty list of SQL arguments.
 - b) The subject routine of RI is defined as follows:
 - i) If RN does not contain a <schema name>, then:
 - 1) Case:
 - A) If RI is immediately contained in a <method selection>, <static method selection>, or a <constructor method selection>, then let DP be TP .
 - B) If the routine execution context of the current SQL-session indicates that an SQL-invoked routine is active, then let DP be the routine SQL-path of that routine execution context.
 - C) Otherwise,
Case:
 - I) 04 If RI is contained in a <schema definition>, then let DP be the SQL-path of that <schema definition>.
 - II) If RI is contained in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then let DP be the SQL-path of the current SQL-session.
 - III) Otherwise, let DP be the SQL-path of the <SQL-client module definition> that contains RI .
 - 2) The *subject routine* of RI is an SQL-invoked routine $SIRSR$ such that:
 - A) $SIRSR$ is an invocable routine.
 - B) The <schema name> of the schema of $SIRSR$ is in DP .
 - C) Case:
 - I) If the routine descriptor of $SIRSR$ does not include a STATIC indication, then there is no other invocable routine $R2$ for which the the <schema name> of the schema that includes $R2$ precedes in DP the <schema name> of the schema that includes $SIRSR$.
 - II) If the routine descriptor of $SIRSR$ includes a STATIC indication, then there is no other invocable routine $R2$ for which the user-defined type described by the descriptor that includes the routine descriptor of $R2$ is a subtype of the user-defined type described by the user-defined type descriptor that includes the routine descriptor of $SIRSR$.
 - ii) If RN contains a <schema name> SN , then SN shall be the <schema name> of a schema S . The *subject routine* of RI is the invocable routine (if any) contained in S .

- c) There shall be exactly one subject routine of *RI*.
- d) If *RI* is not immediately contained in a <call statement>, then the *effective returns data type* of *RI* is the result data type of the subject routine of *RI*.
- e) Let the *static SQL argument list* of *RI* be an empty list of SQL arguments.

8) If <SQL argument list> immediately contains at least one <SQL argument>, then:

- a) The <data type> of each <value expression> immediately contained in a <generalized expression> shall be a subtype of the structured type identified by the <user-defined type name> simply contained in the <path-resolved user-defined type name> that is immediately contained in <generalized expression>.
- b) The set of *candidate routines* of *RI* is defined as follows.

Case:

- i) If *RN* does not contain a <schema name>, then:

1) Case:

- A) If *RI* is immediately contained in a <method selection>, a <static method selection>, or a <constructor method selection>, then let *DP* be *TP*.
- B) If the routine execution context of the current SQL-session indicates that an SQL-invoked routine is active, then let *DP* be the routine SQL-path of that routine execution context.
- C) Otherwise,

Case:

- I) 04 If *RI* is contained in a <schema definition>, then let *DP* be the SQL-path of that <schema definition>.
- II) If *RI* is contained in a <preparable statement> that is prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> or in a <direct SQL statement> that is invoked directly, then let *DP* be the SQL-path of the current SQL-session.
- III) Otherwise, let *DP* be the SQL-path of the <SQL-client module definition> that contains *RI*.

- 2) The candidate routines of *RI* are the set union of invocable routines of all schemas whose <schema name> is in *DP*.

- ii) If *RN* contains a <schema name> *SN*, then *SN* shall be the <schema name> of a schema *S*. The candidate routines of *RI* are the invocable routines (if any) contained in *S*.

c) Case:

- i) If *RI* is immediately contained in a <call statement>, then:

1) Let *XAL* be *AL*.

- 2) The subject routine *SR* of *XAL* is the SQL-invoked routine *SIRCR1* that is a candidate routine of *RI* such that there is no other candidate routine *R2* for which the <schema name> of the schema that includes *R2* precedes in *DP* the <schema name> of the schema that includes *SIRCR1*.

- 3) Let PL be the list of SQL parameters P_i of SR .
 - 4) For each P_i that is an output SQL parameter or both an input SQL parameter and an output SQL parameter, A_i shall be a <target specification>.
 - A) If RI is contained in a <triggered SQL statement> of an AFTER trigger, then A_i shall not be a <column reference>.
 - B) 10 If A_i is an <embedded variable specification> or a <host parameter specification>, then P_i shall be assignable to A_i , according to the Syntax Rules of Subclause 9.1, “Retrieval assignment”, with A_i and P_i as *TARGET* and *VALUE*, respectively.
 - C) 04 If A_i is an <SQL parameter reference>, a <column reference>, or a <target array element specification>, then P_i shall be assignable to A_i , according to the Syntax Rules of Subclause 9.2, “Store assignment”, with A_i and P_i as *TARGET* and *VALUE*, respectively.

NOTE 245 — The <column reference> can only be a new transition variable column reference.
 - 5) For each P_i that is an input SQL parameter but not an output SQL parameter, A_i shall be a <value expression>.
 - 6) For each P_i that is an input SQL parameter or both an input SQL parameter and an output SQL parameter, A_i shall be assignable to P_i , according to the Syntax Rules of Subclause 9.2, “Store assignment”, with P_i and A_i as *TARGET* and *VALUE*, respectively.
- ii) Otherwise:
- 1) A_i shall be a <value expression> or <generalized expression>.
 - 2) Case:
 - A) If A_i is a <generalized expression>, then let TS_i be the data type identified by the <user-defined type name> simply contained in the <path-resolved user-defined type name> that is immediately contained in the <generalized expression>.
 - B) Otherwise, let TS_i be the data type whose data type name is included in the data type descriptor of the data type of A_i .
 - 3) The *subject routine* is defined as follows:
 - A) For each A_i ,

Case:

 - I) If A_i is a <dynamic parameter specification>, then let V_i be A_i .
 - II) Otherwise, let V_i be a value arbitrarily chosen whose declared type is TS_i .
 - B) Let XAL be an <SQL argument list> with N <SQL argument>s derived from the V_i s ordered according to their ordinal position i in XAL .
 - C) The Syntax Rules of Subclause 9.4, “Subject routine determination”, are applied with XAL as *SQL ARGUMENT LIST* and the candidate routines of RI as *SET OF SQL-*

INVOKED ROUTINES, yielding a set of candidate subject routines *CSR* returned from *SET OF SUBJECT ROUTINES*.

D) Case:

I) If *RN* contains a <schema name>, then there shall be exactly one candidate subject routine in *CSR*. The subject routine *SR* is the candidate subject routine in *CSR*.

II) Otherwise:

1) There shall be at least one candidate subject routine in *CSR*.

2) Case:

a) If there is exactly one candidate subject routine in *CSR*, then the subject routine *SR* is the candidate subject routine in *CSR*.

b) If there is more than one candidate subject routine in *CSR*, then

Case:

i) If *RI* is not immediately contained in a <static method selection>, then there shall be an SQL-invoked routine *SIRCR2* in *CSR* such that there is no other candidate subject routine *R2* in *CSR* for which any of the following is true:

1) The <schema name> of the schema that includes *R2* precedes in *DP* the <schema name> of the schema that includes *SIRCR2*.

2) The <schema name> of the schema that includes *R2* is equivalent to the <schema name> of the schema that includes *SIRCR2*.

The subject routine *SR* is *SIRCR2*.

ii) Otherwise, there shall be an SQL-invoked routine *SIRCR3* in *CSR* such that there is no other candidate subject routine *R2* in *CSR* for which the user-defined type described by the user-defined type descriptor that includes the routine descriptor of *R2* is a subtype of the user-defined type described by the user-defined type descriptor that includes the routine descriptor of *SIRCR3*. The subject routine *SR* is *SIRCR3*.

4) The subject routine of *RI* is the subject routine *SR*.

5) Let *PL* be the list of SQL parameters P_i of *SR*.

6) For each P_i , A_i shall be assignable to P_i according to the Syntax Rules of Subclause 9.2, “Store assignment”, with P_i as *TARGET* and A_i as *VALUE*.

7) The *effective returns data type* of *RI* is defined as follows:

A) Case:

I) If *SR* is a type-preserving function, then let P_i be the result SQL parameter of *SR*. If A_i contains a <generalized expression>, then let *RT* be the declared type

of the <value expression> contained in the <generalized expression> of A_i ; otherwise, let RT be the declared type of A_i .

II) Otherwise, let RT be the result data type of SR .

B) The *effective returns data type* of RI is RT .

9) If SR is a constructor function, then RI shall be simply contained in a <new invocation>.

Access Rules

None.

General Rules

1) Let SAL be the *STATIC SQL ARG LIST* and let SR be the *SUBJECT ROUTINE* specified in an application of this Subclause.

NOTE 246 — “static SQL argument list” and “subject routine” are defined by the Syntax Rules of this Subclause.

2) Case:

a) If SAL is empty, then let the *dynamic SQL argument list* DAL be SAL .

b) Otherwise:

i) Each SQL argument A_i in SAL is evaluated, in an implementation-dependent order, to obtain a value V_i .

ii) Let the *dynamic SQL argument list* DAL be the list of values V_i in order.

iii) If SR is type preserving and the null value is substituted for the result parameter, then

Case:

1) If SR is a mutator function, then an exception condition is raised: *data exception — null value substituted for mutator subject parameter* and no further General Rules of this Subclause are applied.

2) Otherwise, the value of RI is the null value and no further General Rules of this Subclause are applied.

iv) Case:

1) If SR is an instance SQL-invoked method, then:

A) If V_1 is the null value, then the value of RI is the null value and the remaining General Rules of this Subclause are not applied.

B) Let SM be the set of SQL-invoked methods M that satisfy the following conditions:

I) The <routine name> of SR and the <routine name> of M have equivalent <qualified identifier>s.

- II) SR and M have the same number N of SQL parameters. Let PSR_i , $1 \text{ (one)} \leq i \leq N$, be the i -th SQL parameter of SR and PM_i , $1 \text{ (one)} \leq i \leq N$, be the i -th SQL parameter of M .
- III) The declared type of the subject parameter of M is a subtype of the declared type of the subject parameter of SR .
- IV) For j varying from 2 to N , the Syntax Rules of [Subclause 9.18](#), “Data type identity”, are applied with the declared type of PM_j and the declared type of PSR_j .

NOTE 247 — SR is an element of the set SM .

- C) SM is the *set of overriding methods* of SR and every SQL-invoked method M in SM is an *overriding method* of SR .

D) Case:

- I) If the first SQL argument AI in SAL contains a <generalized expression>, then let DTI be the data type identified by the <user-defined type name> contained in the <generalized expression> of AI .
- II) Otherwise, let DTI be the most specific type of V_1 .

- E) Let R be the SQL-invoked method in SM such that there is no other SQL-invoked method MI in SM for which the type designator of the declared type of the subject parameter of MI precedes that of the declared type of the subject parameter of R in the type precedence list of DTI .

2) Otherwise, let R be SR .

- 3) Let N and PN be the number of values V_i in DAL . Let T_i be the declared type of the i -th SQL parameter P_i of R . For i ranging from 1 (one) to PN ,

Case:

- a) [14](#) If P_i is an input SQL parameter or both an input SQL parameter and an output SQL parameter, then let CPV_i be the result an application of the General Rules of [Subclause 9.2](#), “Store assignment”, with V_i as *VALUE* and T_i as *TARGET*.

b) Otherwise,

Case:

- i) [13](#) If R is an SQL routine, then let CPV_i be the null value.
- ii) Otherwise, let CPV_i be an implementation-defined value of most specific type T_i .

- 4) [13](#) If R is an external routine, then:

- a) Let P be the program identified by the external name of R .
- b) For i ranging from 1 (one) to N , let P_i be the i -th SQL parameter of R and let T_i be the declared type of P_i .

Case:

- i) If P_i is an input SQL parameter or both an input SQL parameter and an output SQL parameter, then

Case:

- 1) 14 If P_i is a locator parameter, then CPV_i is replaced by the locator value that uniquely identifies the value of CPV_i .
- 2) If T_i is a user-defined type, and P_i is not a locator parameter, then:
 - A) Let FSF_i be the SQL-invoked routine identified by the specific name of the from-sql function associated with P_i in the routine descriptor of R . Let RT_i be the result data type of FSF_i .
 - B) The General Rules of this Subclause are applied with a static SQL argument list that has a single argument that is CPV_i and subject routine FSF_i .
 - C) Let RV_i be the result of the invocation of FSF_i . CPV_i is replaced by RV_i .

- ii) Otherwise,

Case:

- 1) 14 If P_i is a locator parameter, then CPV_i is replaced with an implementation-dependent value of type INTEGER.
- 2) If T_i is a user-defined type and P_i is not a locator parameter, then:
 - A) Let FSF_i be the SQL-invoked routine identified by the specific name of the from-sql function associated with P_i in the routine descriptor of R . Let RT_i be the result data type of FSF_i .
 - B) CPV_i is replaced by an implementation-defined value of type RT_i .

- 5) Preserve the current SQL-session context CSC and create a new SQL-session context RSC derived from CSC as follows:
 - a) Set the current default catalog name, the current default unqualified schema name, the current default character set name, the SQL-path of the current SQL-session, the current default time zone displacement of the current SQL-session, and the contents of all SQL dynamic descriptor areas to implementation-defined values.
 - b) Set the values of the current SQL-session identifier, the identities of all instances of global temporary tables, the current constraint mode for each integrity constraint, the current access mode, the current isolation level, and the current condition area limit to their values in CSC .
 - c) The diagnostics area stack in CSC is copied to RSC and the General Rules of [Subclause 23.2, “Pushing and popping the diagnostics area stack”](#), are applied with “PUSH” as *OPERATION* and the diagnostics area stack in RSC as *STACK*.
 - d) Case:
 - i) 04 If R is an SQL routine, then remove from RSC the identities of all instances of created local temporary tables, declared local temporary tables that are defined by <temporary table declara-

tion>s that are contained in <SQL-client module definition>s, and the cursor instance descriptors of all open cursors that are not global extended dynamic cursors.

- ii) 13 Otherwise:
 - 1) Remove from *RSC* the identities of all instances of created local temporary tables that are referenced in <SQL-client module definition>s that are not the <SQL-client module definition> of *P*, declared local temporary tables that are defined by <temporary table declaration>s that are contained in <SQL-client module definition>s that are not the <SQL-client module definition> of *P*, and the cursor instance descriptors of all open cursors that are not global extended dynamic cursors and whose SQL-client module is not the SQL-client module of *P*.
 - 2) It is implementation-defined whether the identities of all instances of created local temporary tables that are referenced in the <SQL-client module definition> of *P*, declared local temporary tables that are defined by <temporary table declaration>s that are contained in the <SQL-client module definition> of *P*, and the cursor instance descriptors of all open cursors that are not global extended dynamic cursors and SQL-client module is the SQL-client module of *P* are removed from *RSC*.
- e) Indicate in the routine execution context of *RSC* that the SQL-invoked routine *R* is active.
- f) Case:
 - i) If the SQL-data access indication of *CSC* specifies possibly contains SQL and *R* possibly reads SQL-data or *R* possibly modifies SQL-data, then:
 - 1) If *R* is an external routine, then an exception condition is raised: *external routine exception — reading SQL-data not permitted*.
 - 2) Otherwise, an exception condition is raised: *SQL routine exception — reading SQL-data not permitted*.
 - ii) If the SQL-data access indication of *CSC* specifies possibly reads SQL-data and *R* possibly modifies SQL-data, then:
 - 1) If *R* is an external routine, then an exception condition is raised: *external routine exception — modifying SQL-data not permitted*.
 - 2) Otherwise, an exception condition is raised: *SQL routine exception — modifying SQL-data not permitted*.
- g) Case:
 - i) If *R* does not possibly contain SQL, then set the SQL-data access indication in the routine execution context of *RSC* to *does not possibly contain SQL*.
 - ii) If *R* possibly contains SQL, then set the SQL-data access indication in the routine execution context of *RSC* to *possibly contains SQL*.
 - iii) If *R* possibly reads SQL-data, then set the SQL-data access indication in the routine execution context of *RSC* to *possibly reads SQL-data*.
 - iv) If *R* possibly modifies SQL-data, then set the SQL-data access indication in the routine execution context of *RSC* to *possibly modifies SQL-data*.
- h) The authorization stack of *RSC* is set to a copy of the authorization stack of *CSC*.

- i) A copy of the top cell is pushed onto the authorization stack of *RSC*.
 - j) Case:
 - i) If *R* is an external routine, then:
 - 1) Case:
 - A) If the external security characteristic of *R* is IMPLEMENTATION DEFINED, then the current user identifier and the current role name of *RSC* are implementation-defined.
 - B) If the external security characteristic of *R* is DEFINER, then the top cell of the authorization stack of *RSC* is set to contain only the routine authorization identifier of *R*.
 - 2) The routine SQL-path of *RSC* is set to the external routine SQL-path of *R*.
 - ii) Otherwise:
 - 1) If the SQL security characteristic of *R* is DEFINER, then the top cell of the authorization stack of *RSC* is set to contain only the routine authorization identifier of *R*.
 - 2) The routine SQL-path of *RSC* is set to the routine SQL-path of *R*.
 - k) If the subject routine is an SQL-invoked procedure *SIP*, then:
 - i) For every received cursor *RC* whose origin is the <specific routine designator> of *SIP*, the cursor declaration descriptor and the cursor instance descriptor of *RC* are destroyed.
 - ii) Every result set sequence associated with *SIP* is destroyed.
 - iii) Let *INV* be the invoker of *SIP*. An empty result set sequence *RSS*, for SQL-invoked procedure *SIP* and invoker *INV*, is added to *RSC*.

NOTE 248 — The invoker of *SIP* is defined in Subclause 4.27.5, “Result sets returned by SQL-invoked procedures”.
 - l) *RSC* becomes the current SQL-session context.
- 6) If the descriptor of *R* includes an indication that a new savepoint level is to be established when *R* is invoked, then a new savepoint level is established.
- 7) If *R* is an SQL routine, then
- Case:
- a) If *R* is a null-call function and if any of *CPV_i* is the null value, then let *RV* be the null value.
 - b) Otherwise:
 - i) For *i* ranging from 1 (one) to *PN*, set the value of *P_i* to *CPV_i*.
 - ii) The General Rules of Subclause 13.5, “<SQL procedure statement>”, are evaluated with the SQL routine body of *R* as the *executing statement*.
 - iii) If, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-connection statement, then an exception condition is raised: *SQL routine exception — prohibited SQL-statement attempted*.
 - iv) Case:

- 1) If the SQL-implementation does not support Feature T272, “Enhanced savepoint management”, and, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-transaction statement, then an exception condition is raised: *SQL routine exception — prohibited SQL-statement attempted*.
 - 2) If, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-transaction statement that is not a <savepoint statement> or a <release savepoint statement>, or is a <rollback statement> that does not specify a <savepoint clause>, then an exception condition is raised: *SQL routine exception — prohibited SQL-statement attempted*.
 - v) If the SQL implementation does not support Feature T651, “SQL-schema statements in SQL routines”, and, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-schema statement, an exception condition is raised: *SQL routine exception — prohibited SQL-statement attempted*.
 - vi) If the SQL implementation does not support Feature T652, “SQL-dynamic statements in SQL routines”, and, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-dynamic statement, an exception condition is raised: *SQL routine exception — prohibited SQL-statement attempted*.
 - vii) If the SQL-data access indication of *RSC* specifies *possibly contains SQL* and, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-statement that possibly reads SQL-data, or an attempt is made to execute an SQL-statement that possibly modifies SQL-data, then an exception condition is raised: *SQL routine exception — reading SQL-data not permitted*.
 - viii) If the SQL-data access indication of *RSC* specifies *possibly reads SQL-data* and, before the completion of the execution of the SQL routine body of *R*, an attempt is made to execute an SQL-statement that possibly modifies SQL-data then an exception condition is raised: *SQL routine exception — modifying SQL-data not permitted*.
 - ix) If *R* is an SQL-invoked function, then
Case:
 - 1) If no <return statement> is executed before completion of the execution of the SQL routine body of *R*, then an exception condition is raised: *SQL routine exception — function executed no return statement*.
 - 2) Otherwise, let *RV* be the returned value of the execution of the SQL routine body of *R*.
NOTE 249 — “Returned value” is defined in Subclause 16.2, “<return statement>”.
 - x) If *R* is an SQL-invoked procedure, then for each SQL parameter of *R* that is an output SQL parameter or both an input SQL parameter and an output SQL parameter, set the value of *CPV_i* to the value of *P_i*.
- 8) If *R* is an external routine, then:
- a) 04 The method and time of binding of *P* to the schema that includes *R* is implementation-defined.
 - b) If *R* specifies PARAMETER STYLE SQL, then
 - i) Case:

- 1) If R is an SQL-invoked function, then the effective SQL parameter list $ESPL$ of R is set as follows:
 - A) If R is an array-returning external function or a multiset-returning external function with the element type being a row type, then let FRN be the degree of the element type; otherwise, let FRN be 1 (one).
 - B) For i ranging from 1 (one) to PN , the i -th entry in $ESPL$ is set to CPV_i .
 - C) For i ranging from $PN+1$ to $PN+FRN$, the i -th entries in $ESPL$ are the *result data items*.
 - D) For i ranging from $(PN+FRN)+1$ to $(PN+FRN)+N$, the i -th entry in $ESPL$ is the *SQL indicator argument* corresponding to $CPV_{i-(PN+FRN)}$.
 - E) For i ranging from $(PN+FRN)+N+1$ to $(PN+FRN)+N+FRN$, the i -th entries in $ESPL$ are the SQL indicator arguments corresponding to the result data items.
 - F) For i equal to $(PN+FRN)+(N+FRN)+1$, the i -th entry in $ESPL$ is the *exception data item*.
 - G) For i equal to $(PN+FRN)+(N+FRN)+2$, the i -th entry in $ESPL$ is the *routine name text item*.
 - H) For i equal to $(PN+FRN)+(N+FRN)+3$, the i -th entry in $ESPL$ is the *specific name text item*.
 - I) For i equal to $(PN+FRN)+(N+FRN)+4$, the i -th entry in $ESPL$ is the *message text item*.
 - J) If R is an array-returning external function or a multiset-returning external function, then for i equal to $(PN+FRN)+(N+FRN)+5$, the i -th entry in $ESPL$ is the *save area data item* and for i equal to $(PN+FRN)+(N+FRN)+6$, the i -th entry in $ESPL$ is the *call type data item*.
 - K) Set the values of the SQL indicator arguments corresponding to the result data items (that is, SQL argument value list entries from $(PN+FRN)+N+1$ through $(PN+FRN)+N+FRN$, inclusive) to 0 (zero).
 - L) For i ranging from 1 (one) to PN , if CPV_i is the null value, then set entry $(PN+FRN)+i$ (that is, the i -th SQL indicator argument corresponding to CPV_i) to -1; otherwise, set entry $(PN+FRN)+i$ (that is, the i -th SQL indicator argument corresponding to CPV_i) to 0 (zero).
 - M) If R is an array-returning external function or a multiset-returning external function, then set the value of the save area data item (that is, SQL argument value list entry $(PN+FRN)+(N+FRN)+5$) to 0 (zero) and set the value of the call type data item (that is, SQL argument value list entry $(PN+FRN)+(N+FRN)+6$) to -1.
- 2) Otherwise, the effective SQL parameter list $ESPL$ of R is set as follows:
 - A) For i ranging from 1 (one) to PN , the i -th entry in $ESPL$ is CPV_i .
 - B) For i ranging from $PN+1$ to $PN+N$, the i -th entry in $ESPL$ is the *SQL indicator argument* corresponding to CPV_{i-PN} .
 - C) For i equal to $(PN+N)+1$, the i -th entry in $ESPL$ is the *exception data item*.
 - D) For i equal to $(PN+N)+2$, the i -th entry in $ESPL$ is the *routine name text item*.

- E) For i equal to $(PN+N)+3$, the i -th entry in *ESPL* is the *specific name text item*.
 - F) For i equal to $(PN+N)+4$, the i -th entry in *ESPL* is the *message text item*.
 - G) For i ranging from 1 (one) to PN , if CPV_i is the null value, then set entry $PN+i$ in *ESPL* (that is, the i -th SQL indicator argument corresponding to CPV_i) to -1 ; otherwise, set entry $PN+i$ in *ESPL* (that is, the i -th SQL indicator argument corresponding to CPV_i) to 0 (zero).
- ii) The exception data item is set to '00000'.
 - iii) The routine name text item is set to the <schema qualified name> of the routine name of R .
 - iv) The specific name text item is set to the <qualified identifier> of the specific name of R .
 - v) The message text item is set to a zero-length string.
- c) If R specifies PARAMETER STYLE GENERAL, then the effective SQL parameter list *ESPL* of R is set as follows:
- i) If R is not a null-call function and, for i ranging from 1 (one) to PN , CPV_i is the null value, then an exception condition is raised: *external routine invocation exception — null value not allowed*.
 - ii) For i ranging from 1 (one) to PN , if no CPV_i is the null value, then for j ranging from 1 (one) to PN , the j -th entry in *ESPL* is set to CPV_j .
- d) 13 If R specifies DETERMINISTIC and if different executions of P with identical SQL argument value lists do not produce identical results, then the results are implementation-dependent.
- e) Let EN be the number of entries in *ESPL*. Let ESP_i be the i -th effective SQL parameter in *ESPL*.
- f) Case:
- i) If R is a null-call function and if any of CPV_i is the null value, then P is assumed to have been executed.
 - ii) Otherwise:
 - 1) 13 If R is not an array-returning external function or a multiset-returning external function, then P is executed with a list of EN parameters PD_i whose parameter names are PN_i and whose values are set as follows:
 - A) If the language of R is ADA, C, COBOL, FORTRAN, M, PASCAL, or PLI, then let the *operative data type correspondences table* be Table 16, “Data type correspondences for Ada”, Table 17, “Data type correspondences for C”, Table 18, “Data type correspondences for COBOL”, Table 19, “Data type correspondences for Fortran”, Table 20, “Data type correspondences for M”, Table 21, “Data type correspondences for Pascal”, or Table 22, “Data type correspondences for PL/I”, respectively. Refer to the two columns of the operative data type correspondences table as the “SQL data type” column and the “host data type” column.
 - B) For i varying from 1 (one) to EN , the data type DT_i of PD_i is the data type listed in the host data type column of the row in the data type correspondences table whose value in the SQL data type column corresponds to the data type of ESP_i .

- C) The value of PD_i is set to the value of ESP_i .
- 2) 13 If R is an array-returning external function, then:
- A) Let AR be an array whose declared type is the result data type of R .
 - B) The General Rules of Subclause 9.14, “Execution of array-returning functions”, are applied with AR as *ARRAY*, $ESPL$ as *EFFECTIVE SQL PARAMETER LIST*, and P as *PROGRAM*.
- 3) If R is a multiset-returning external function, then:
- A) Let MU be a multiset whose declared type is the result data type of R .
 - B) The General Rules of Subclause 9.15, “Execution of multiset-returning functions”, are applied with MU as *MULTISET*, $ESPL$ as *EFFECTIVE SQL PARAMETER LIST*, and P as *PROGRAM*.
- 4) 13 If the SQL-data access indication of RSC specifies *does not possibly contain SQL* and, before the completion of any execution of P , an attempt is made to execute an SQL-statement, then an exception condition is raised: *external routine exception — containing SQL not permitted*.
- 5) If, before the completion of any execution of P , an attempt is made to execute an SQL-connection statement, then an exception condition is raised: *external routine exception — prohibited SQL-statement attempted*.
- 6) Case:
- A) If the SQL-implementation does not support Feature T272, “Enhanced savepoint management”, and, before the completion of the execution of P , an attempt is made to execute an SQL-transaction statement, then an exception condition is raised: *external routine exception — prohibited SQL-statement attempted*.
 - B) If, before the completion of the execution of P , an attempt is made to execute an SQL-transaction statement that is not <savepoint statement> or <release savepoint statement>, or is a <rollback statement> that does not specify a <savepoint clause>, then an exception condition is raised: *external routine exception — prohibited SQL-statement attempted*.
- 7) If the SQL implementation does not support Feature T653, “SQL-schema statements in external routines”, and, before the completion of any execution of P , an attempt is made to execute an SQL-schema statement, then an exception condition is raised: *external routine exception — prohibited SQL-statement attempted*.
- 8) If the SQL implementation does not support Feature T654, “SQL-dynamic statements in external routines”, and, before the completion of any execution of P , an attempt is made to execute an SQL-dynamic statement, then an exception condition is raised: *external routine exception — prohibited SQL-statement attempted*.
- 9) If the SQL-data access indication of RSC specifies *possibly contains SQL* and, before the completion of any execution of P , an attempt is made to execute an SQL-statement that possibly reads SQL-data, or an attempt is made to execute an SQL-statement that possibly modifies SQL-data, then an exception condition is raised: *external routine exception — reading SQL-data not permitted*.

- 10) If the SQL-data access indication of *RSC* specifies *possibly reads SQL-data* and, before the completion of any execution of *P*, an attempt is made to execute an SQL-statement that possibly modifies SQL-data, then an exception condition is raised: *external routine exception — modifying SQL-data not permitted*.
 - 11) If *P* is not a conforming program according to the standard for *R*, then the results of any execution of *P* are implementation-dependent.
- g) After the completion of any execution of *P*:
- i) 13 It is implementation-defined whether:
 - 1) Every instance of created local temporary tables and every instance of declared local temporary tables that is associated with *RSC* is destroyed.
 - 2) For every prepared statement *PS* prepared by *P* in the current SQL-transaction that has not been deallocated by *P*:
 - A) Let *SSN* be the <SQL statement name> that identifies *PS*.
 - B) The following SQL-statement is effectively executed:


```
DEALLOCATE PREPARE SSN
```
 - ii) 13 For *i* varying from 1 (one) to *EN*, the value of *ESP_i* is set to the value of *PD_i*. If *R* specifies PARAMETER STYLE SQL, then

Case:

 - 1) If the exception data item has the value '00000', then the execution of *P* was successful.
 - 2) If the first two characters of the exception data item are equal to the SQLSTATE condition code class value for *warning*, then a completion condition is raised: *warning*, using a subclass code equal to the final three characters of the value of the exception data item.
 - 3) Otherwise, an exception condition is raised using a class code equal to the first two characters of the value of the exception data item and a subclass code equal to the final three characters of the value of the exception data item.
 - iii) If the exception data item is not '00000' and *R* specified PARAMETER STYLE SQL, then the message text item is stored in the first diagnostics area.
- h) If *R* is an SQL-invoked function, then:
- i) Case:
 - 1) If *R* is an SQL-invoked method whose routine descriptor does not include a STATIC indication and if *CPV₁* is the null value, then let *RDI* be the null value.
 - 2) If *R* is a null-call function, *R* is not an array-returning external function or a multiset-returning external function, and if any of *CPV_i* is the null value, then let *RDI* be the null value.
 - 3) 13 If *R* is not a null-call function, *R* is not an array-returning function or a multiset-returning function, *R* specifies PARAMETER STYLE SQL, and entry $(PN+1)+N+1$ in *ESPL* (that is, SQL indicator argument *N*+1 corresponding to the result data item) is negative, then let *RDI* be the null value.

4) Otherwise,

A) Case:

- I) If *R* is not an array-returning external function or a multiset-returning external function, *R* specifies PARAMETER STYLE SQL, and entry $(PN+1)+N+1$ in *ESPL* (that is, SQL indicator argument *N*+1 corresponding to the result data item) is not negative, then let *ERDI* be the value of the result data item.
- II) If *R* is an array-returning external function, and *R* specifies PARAMETER STYLE SQL, then let *ERDI* be *AR*.
- III) If *R* is a multiset-returning function, and *R* specifies PARAMETER STYLE SQL, then let *ERDI* be *MU*.
- IV) If *R* specifies PARAMETER STYLE GENERAL, then let *ERDI* be the value returned from *P*.

NOTE 250 — The value returned from *P* is passed to the SQL-implementation in an implementation-dependent manner. An argument value list entry is not used for this purpose.

B) Case:

- I) If *ERDI* is an invalid locator, then an exception condition is raised: *locator exception — invalid specification* and no further General Rules of this Subclause are applied.
- II) If the routine descriptor of *R* indicates that the return value is a locator, then

Case:

- 1) If *RT* is a binary large object type, then let *RDI* be the binary large object value corresponding to *ERDI*.
- 2) If *RT* is a character large object type, then let *RDI* be the large object character string corresponding to *ERDI*.
- 3) If *RT* is an array type, then let *RDI* be the array value corresponding to *ERDI*.
- 4) If *RT* is a multiset type, then let *RDI* be the multiset value corresponding to *ERDI*.
- 5) If *RT* is a user-defined type, then let *RDI* be the user-defined type value corresponding to *ERDI*.
- III) Otherwise, if *R* specifies <result cast>, then let *CRT* be the <data type> specified in <result cast>; otherwise, let *CRT* be the <returns data type> of *R*.

Case:

- 1) If *R* specifies <result cast> and the routine descriptor of *R* indicates that the <result cast> has a locator indication, then

Case:

- a) If *CRT* is a binary large object type, then let *RDI* be the binary large object value corresponding to *ERDI*.

- b) If *CRT* is a character large object type, then let *RDI* be the large object character string corresponding to *ERDI*.
- c) If *CRT* is an array type, then let *RDI* be the array value corresponding to *ERDI*.
- d) If *CRT* is a multiset type, then let *RDI* be the multiset value corresponding to *ERDI*.
- e) If *CRT* is a user-defined type, then let *RDI* be the user-defined type value corresponding to *ERDI*.

2) Otherwise,

Case:

a) If *CRT* is a user-defined type, then:

- i) Let *TSF* be the SQL-invoked routine identified by the specific name of the to-sql function associated with the result of *R*.
- ii) Case:
 - 1) If *TSF* is an SQL-invoked method, then:
 - A) If *R* is a type-preserving function, then let *MAT* be the most specific type of the value of the argument substituted for the result SQL parameter of *R*; otherwise, let *MAT* be *CRT*.
 - B) The General Rules of this Subclause are applied with a static SQL argument list whose first element is the value returned by the invocation of:

MAT ()

and whose second element is *ERDI*, and the subject routine *TSF*.

- 2) Otherwise, the General Rules of this Subclause are applied with a static SQL argument list that has a single SQL argument that is *ERDI*, and the subject routine *TSF*.

iii) Let *RDI* be the result of invocation of *TSF*.

b) 14 Otherwise, let *RDI* be *ERDI*.

- ii) If *R* specified a <result cast>, then let *RT* be the <returns data type> of *R* and let *RV* be the result of:

CAST (*RDI* AS *RT*)

Otherwise, let *RV* be *RDI*.

- i) If *R* is an SQL-invoked procedure, then for each *P_i*, 1 (one) ≤ *i* ≤ *PN*, that is an output SQL parameter or both an input SQL parameter and an output SQL parameter,

Case:

- i) If R specifies PARAMETER STYLE SQL and entry $(PN+i)$ in $ESPL$ (that is, the i -th SQL indicator argument corresponding to CPV_i) is negative, then CPV_i is set to the null value.
- ii) 13 If R specifies PARAMETER STYLE SQL, and entry $(PN+i)$ in $ESPL$ (that is, the i -th SQL indicator argument corresponding to CPV_i) is not negative, and a value was not assigned to the i -th entry in $ESPL$, then CPV_i is set to an implementation-defined value of type T_i .
- iii) Otherwise:

NOTE 251 — In this case, either R specifies PARAMETER STYLE SQL and entry $(PN+1)+i$ in $SQPL$ (that is, the i -th SQL indicator argument corresponding to CPV_i) is not negative and a value was assigned to the i -th entry in $ESPL$, or else R specifies PARAMETER STYLE GENERAL.

- 1) Let EV_i be the i -th entry in $ESPL$. Let T_i be the <data type> of P_i .
- 2) If EV_i is an invalid locator, then an exception condition is raised: *locator exception — invalid specification* and no further General Rules of this Subclause are applied.
- 3) Case:

A) If P_i is a locator parameter, then

Case:

- I) If T_i is a binary large object type, then CPV_i is set to the binary large object value corresponding to EV_i .
- II) If T_i is a character large object type, then CPV_i is set to the large object character string corresponding to EV_i .
- III) If T_i is an array type, then CPV_i is set to the array value corresponding to EV_i .
- IV) If T_i is a multiset type, then CPV_i is set to the multiset value corresponding to EV_i .
- V) If T_i is a user-defined type, then CPV_i is set to the user-defined type value corresponding to EV_i .

B) If T_i is a user-defined type, then:

- I) Let TSF_i be the SQL-invoked function identified by the specific name of the to-sql function associated with P_i in the routine descriptor of R .
- II) Case:
 - 1) If TSF_i is an SQL-invoked method, then the General Rules of this Subclause are applied with a static SQL argument list whose first element is the value returned by the invocation of:

$T_i()$

and whose second element is EV_i , and the subject routine TSF_i .

- 2) Otherwise, the General Rules of this Subclause are applied with a static SQL argument list that has a single SQL argument that is EV_i , and the subject routine TSF_i .

III) CPV_i is set to the result of an invocation of TSF_i .

- C) 14 Otherwise, CPV_i is set to EV_i .

9) Case:

- a) If R is an SQL-invoked function, then:

- i) If R is a type-preserving function, then:

- 1) Let MAT be the most specific type of the value of the argument substituted for the result SQL parameter of R .
- 2) If RV is not the null value and the most specific type of RV is not compatible with MAT , then an exception condition is raised: *data exception — most specific type mismatch*.

- ii) Let $ERDT$ be the effective returns data type of the <routine invocation>.

- iii) 14 Let the result of the <routine invocation> be the result of the application of the General Rules of Subclause 9.2, “Store assignment”, with RV as *VALUE* and $ERDT$ as *TARGET*.

- b) Otherwise, for each SQL parameter P_i of R that is an output SQL parameter or both an input SQL parameter and an output SQL parameter, let TS_i be the <target specification> of the corresponding <SQL argument> A_i .

Case:

- i) If TS_i is a <host parameter specification> or an <embedded variable specification>, then the General Rules of Subclause 9.1, “Retrieval assignment”, are applied with CPV_i as *VALUE* and TS_i as *TARGET*.

- ii) 04 If TS_i is an <SQL parameter reference>, a <column reference>, or a <target array element specification>, then

NOTE 252 — The <column reference> can only be a new transition variable column reference.

Case:

- 1) If <target array element specification> is specified, then

Case:

- A) If the value of TS_i , denoted by C , is the null value, then an exception condition is raised: *data exception — null value in array target*.

- B) Otherwise:

- I) Let N be the maximum cardinality of C .

- II) Let M be the cardinality of the value of C .

- III) Let I be the value of the <simple value specification> immediately contained in TS_i .

- IV) Let *EDT* be the element type of *C*.
- V) Case:
 - 1) If *I* is greater than zero and less than or equal to *M*, then the value of *C* is replaced by an array *A* with element type *EDT* and cardinality *M* derived as follows:
 - a) For *j* varying from 1 (one) to *I*–1 and from *I*+1 to *M*, the *j*-th element in *A* is the value of the *j*-th element in *C*.
 - b) 14 The *I*-th element of *A* is set to the value of *CPV_i*, denoted by *SV*, by applying the General Rules of Subclause 9.2, “Store assignment”, to the *I*-th element of *A* and *SV* as *TARGET* and *VALUE*, respectively.
 - 2) If *I* is greater than *M* and less than or equal to *N*, then the value of *C* is replaced by an array *A* with element type *EDT* and cardinality *I* derived as follows:
 - a) For *j* varying from 1 (one) to *M*, the *j*-th element in *A* is the value of the *j*-th element in *C*.
 - b) For *j* varying from *M*+1 to *I*, the *j*-th element in *A* is the null value.
 - c) 14 The *I*-th element of *A* is set to the value of *CPV_i*, denoted by *SV*, by applying the General Rules of Subclause 9.2, “Store assignment”, to the *I*-th element of *A* and *SV* as *TARGET* and *VALUE*, respectively.
 - 3) Otherwise, an exception condition is raised: *data exception — array element error*.

- 2) 14 Otherwise, the General Rules of Subclause 9.2, “Store assignment” are applied with *CPV_i* as *VALUE* and *TS_i* as *TARGET*.

10) If *RSS* is not empty, then let *PR* be the descriptor of *SIP*.

- a) Let *MAX* be the maximum number of returned result sets included in *PR*.
- b) 13 Let *OPN* be the actual number of returned result sets included in *RSS*.
- c) Case:
 - i) If *OPN* is greater than *MAX*, then:
 - 1) Let *RTN* be *MAX*.
 - 2) A completion condition is raised: *warning — attempt to return too many result sets*.
 - ii) Otherwise, let *RTN* be *OPN*.
- d) 13 For each *i*, $1 \text{ (one)} \leq i \leq RTN$, let *FRC_i* be the with-return cursor of the *i*-th returned result set *RS_i* in *RSS*.
- e) For each *i*, $1 \text{ (one)} \leq i \leq RTN$,
Case:

i) If FRC_i is a scrollable cursor, then the initial cursor position of RS_i is the current cursor position of FRC_i .

ii) Otherwise,

Case:

1) If the application of Subclause 15.3, “Determination of the current row of a cursor”, with FRC_i as *CURSOR* and NEXT as *FETCH ORIENTATION*, would position the cursor on or before some row in RS_i , then let RN be the ordinal position of that row in RS_i .

2) Otherwise, let RN be one greater than the number of rows in RS_i .

The first RN rows are deleted from RS_i and the initial cursor position of RS_i is before the first row.

f) 13 A completion condition is raised: *warning — result sets returned*.

11) If R is an SQL routine, then for every open cursor CR that was opened during the execution of R and that is either a declared cursor or a local extended dynamic cursor, the General Rules of Subclause 15.4, “Effect of closing a cursor”, are applied, with CR as *CURSOR* and SAVE as *DISPOSITION*.

12) If R is an external routine, then it is implementation-defined whether, for every open cursor CR that was opened during the execution of R and that is either a declared cursor or a local extended dynamic cursor:

a) The the General Rules of Subclause 15.4, “Effect of closing a cursor”, are applied, with CR as *CURSOR* and SAVE as *DISPOSITION*.

b) The cursor instance descriptor of CR is destroyed.

13) 13 Prepare CSC to become the current SQL-session context:

a) Set the value of the current constraint mode for each integrity constraint in CSC to the value of the current constraint mode for each integrity constraint in RSC .

b) Set the value of the current access mode in CSC to the value of the current access mode in RSC .

c) Set the value of the current isolation level in CSC to the value of the current isolation level in RSC .

d) Set the value of the current condition area limit in CSC to the value of the current condition area limit CAL in RSC .

e) For each occupied condition area CA in the first diagnostics area of RSC , if the value of RETURNED_SQLSTATE in CA does not represent *successful completion*, then

Case:

i) If the number of occupied condition areas in the first diagnostics area DAI in CSC is less than CAL , then CA is copied to the first vacant condition area in DAI .

NOTE 253 — This causes the first vacant condition area in DAI to become occupied.

ii) Otherwise, the value of MORE in the statement information area of DAI is set to 'Y'.

f) Replace the identities of all instances of global temporary tables in CSC with the identities of the instances of global temporary tables in RSC .

- g) Remove the top cell from the authorization stack of *RSC* and set the authorization stack of *CSC* to a copy of the authorization stack of *RSC*.

NOTE 254 — The copying of *RSC*'s authorization stack into *CSC* is necessary in order to carry back any change in the SQL-session user identifier.

- h) If the subject routine is an SQL-invoked procedure, then the result set sequence *RSS* is added to *CSC*.

NOTE 255 — *RSS* is now available for reference by an <allocate cursor statement> containing FOR PROCEDURE.

- 14) If *R* is an SQL-invoked function or if *R* is an SQL-invoked procedure and the descriptor of *R* includes an indication that a new savepoint level is to be established when *R* is invoked, then the current savepoint level is destroyed.
- 15) *CSC* becomes the current SQL-session context.

Conformance Rules

- 1) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <generalized expression>.
- 2) Without Feature S201, “SQL-invoked routines on arrays”, conforming SQL language shall not contain an <SQL argument> whose declared type is an array type.
- 3) Without Feature S202, “SQL-invoked routines on multisets”, conforming SQL language shall not contain an <SQL argument> whose declared type is a multiset type.
- 4) Without Feature B033, “Untyped SQL-invoked function arguments”, conforming SQL language shall not contain a <routine invocation> that is not simply contained in a <call statement> that simply contains an <SQL argument> that is a <dynamic parameter specification>.

10.5 <character set specification>

Function

Identify a character set.

Format

```
<character set specification> ::=  
    <standard character set name>  
    | <implementation-defined character set name>  
    | <user-defined character set name>  
  
<standard character set name> ::=  
    <character set name>  
  
<implementation-defined character set name> ::=  
    <character set name>  
  
<user-defined character set name> ::=  
    <character set name>
```

Syntax Rules

- 1) The <standard character set name>s and <implementation-defined character set name>s that are supported are implementation-defined.
- 2) A character set identified by a <standard character set name>, or by an <implementation-defined character set name> has associated with it a privilege descriptor that was effectively defined by the <grant statement>

```
GRANT USAGE ON  
CHARACTER SET CS  
TO PUBLIC
```

where *CS* is the <character set name> contained in the <character set specification>. The grantor of the privilege descriptor is set to the special grantor value “_SYSTEM”.

- 3) The <standard character set name>s shall include SQL_CHARACTER and those character sets specified in [Subclause 4.2.7, “Character sets”](#), as defined by this and other standards.
- 4) The <implementation-defined character set name>s shall include SQL_TEXT and SQL_IDENTIFIER.
- 5) Let *C* be the <character set name> contained in the <character set specification>. The schema identified by the explicit or implicit qualifier of the <character set name> shall include the descriptor of *C*.
- 6) If a <character set specification> is not contained in a <schema definition>, then the <character set name> immediately contained in the <character set definition> shall contain an explicit <schema name> that is not equivalent to INFORMATION_SCHEMA.

Access Rules

- 1) Case:

- a) If <character set specification> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include USAGE on *C*.
- b) Otherwise, the current privileges shall include USAGE on *C*.

General Rules

- 1) A <character set specification> identifies a character set. Let the identified character set be *CS*.
- 2) A <standard character set name> specifies the name of a character set that is defined by a national or international standard. The character repertoire of *CS* is defined by the standard defining the character set identified by that <standard character set name>. The default collation of the character set is defined by the order of the characters in the standard and has the PAD SPACE characteristic.
- 3) An <implementation-defined character set name> specifies the name of a character set that is implementation-defined. The character repertoire of *CS* is implementation-defined. The default collation of the character set and whether the collation has the NO PAD characteristic or the PAD SPACE characteristic is implementation-defined.
- 4) A <user-defined character set name> identifies a character set whose descriptor is included in some schema whose <schema name> is not equivalent to INFORMATION_SCHEMA.

NOTE 256 — The default collation of the character set is defined as in [Subclause 11.32](#), “<character set definition>”.

- 5) There is a character set descriptor for every character set that can be specified by a <character set specification>.

Conformance Rules

- 1) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <character set specification>.

10.6 <specific routine designator>

Function

Specify an SQL-invoked routine.

Format

```
<specific routine designator> ::=  
    SPECIFIC <routine type> <specific name>  
    | <routine type> <member name> [ FOR <schema-resolved user-defined type name> ]  
  
<routine type> ::=  
    ROUTINE  
    | FUNCTION  
    | PROCEDURE  
    | [ INSTANCE | STATIC | CONSTRUCTOR ] METHOD  
  
<member name> ::=  
    <member name alternatives> [ <data type list> ]  
  
<member name alternatives> ::=  
    <schema qualified routine name>  
    | <method name>  
  
<data type list> ::=  
    <left paren> [ <data type> [ { <comma> <data type> }... ] ] <right paren>
```

Syntax Rules

- 1) If a <specific name> *SN* is specified, then the <specific routine designator> shall identify an SQL-invoked routine whose <specific name> is *SN*.
- 2) If <routine type> specifies METHOD and none of INSTANCE, STATIC, or CONSTRUCTOR is specified, then INSTANCE is implicit.
- 3) If a <member name> *MN* is specified, then:
 - a) If <schema-resolved user-defined type name> is specified, then <routine type> shall specify METHOD. If METHOD is specified, then <schema-resolved user-defined type name> shall be specified.
 - b) Case:
 - i) If <routine type> specifies METHOD, then <method name> shall be specified. Let *SCN* be the implicit or explicit <schema name> of <schema-resolved user-defined type name>, let *METH* be the <method name>, and let *RN* be *SCN.METH*.
 - ii) Otherwise, <schema qualified routine name> shall be specified. Let *RN* be the <schema qualified routine name> of *MN* and let *SCN* be the <schema name> of *MN*.
 - c) Case:
 - i) If *MN* contains a <data type list>, then:

10.6 <specific routine designator>

- 1) If <routine type> specifies FUNCTION, then there shall be exactly one SQL-invoked regular function in the schema identified by *SCN* whose <schema qualified routine name> is *RN* such that for all *i* the Syntax Rules of Subclause 9.18, “Data type identity”, when applied with the declared type of its *i*-th SQL parameter and the *i*-th <data type> in the <data type list> of *MN*, are satisfied. The <specific routine designator> identifies that SQL-invoked function.
- 2) If <routine type> specifies PROCEDURE, then there shall be exactly one SQL-invoked procedure in the schema identified by *SCN* whose <schema qualified routine name> is *RN* such that for all *i* the Syntax Rules of Subclause 9.18, “Data type identity”, when applied with the declared type of its *i*-th SQL parameter and the *i*-th <data type> in the <data type list> of *MN*, are satisfied. The <specific routine designator> identifies that SQL-invoked procedure.
- 3) If <routine type> specifies METHOD, then

Case:

- A) If STATIC is specified, then there shall be exactly one static SQL-invoked method of the type identified by <schema-resolved user-defined type name> whose <method name> is *METH*, such that for all *i* the Syntax Rules of Subclause 9.18, “Data type identity”, when applied with the declared type of its *i*-th SQL parameter and the *i*-th <data type> in the <data type list> of *MN*, are satisfied. The <specific routine designator> identifies that static SQL-invoked method.
- B) If CONSTRUCTOR is specified, then there shall be exactly one SQL-invoked constructor method of the type identified by <schema-resolved user-defined type name> whose <method name> is *METH*, such that for all *i* the Syntax Rules of Subclause 9.18, “Data type identity”, when applied with the declared type of its *i*-th SQL parameter in the unaugmented <SQL parameter declaration list> and the *i*-th <data type> in the <data type list> of *MN*, are satisfied. The <specific routine designator> identifies that SQL-invoked constructor method.
- C) Otherwise, there shall be exactly one instance SQL-invoked method of the type identified by <schema-resolved user-defined type name> whose <method name> is *METH*, such that for all *i* the Syntax Rules of Subclause 9.18, “Data type identity”, when applied with the declared type of its *i*-th SQL parameter in the unaugmented <SQL parameter declaration list> and the *i*-th <data type> in the <data type list> of *MN*, are satisfied. The <specific routine designator> identifies that instance SQL-invoked method.
- 4) If <routine type> specifies ROUTINE, then there shall be exactly one SQL-invoked routine in the schema identified by *SCN* whose <schema qualified routine name> is *RN* such that for all *i* the Syntax Rules of Subclause 9.18, “Data type identity”, when applied with the declared type of its *i*-th SQL parameter and the *i*-th <data type> in the <data type list> of *MN*, are satisfied. The <specific routine designator> identifies that SQL-invoked routine.

ii) Otherwise:

- 1) If <routine type> specifies FUNCTION, then there shall be exactly one SQL-invoked function in the schema identified by *SCN* whose <schema qualified routine name> is *RN*. The <specific routine designator> identifies that SQL-invoked function.

- 2) If <routine type> specifies PROCEDURE, then there shall be exactly one SQL-invoked procedure in the schema identified by *SCN* whose <schema qualified routine name> is *RN*. The <specific routine designator> identifies that SQL-invoked procedure.

- 3) If <routine type> specifies METHOD, then

Case:

- A) If STATIC is specified, then there shall be exactly one static SQL-invoked method of the user-defined type identified by <schema-resolved user-defined type name> whose <method name> is *METH*. The <specific routine designator> identifies that static SQL-invoked method.
- B) If CONSTRUCTOR is specified, then there shall be exactly one SQL-invoked constructor method of the user-defined type identified by <schema-resolved user-defined type name> whose <method name> is *METH*. The <specific routine designator> identifies that SQL-invoked constructor method.
- C) Otherwise, there shall be exactly one instance SQL-invoked method of the user-defined type identified by <schema-resolved user-defined type name> whose <method name> is *METH*. The <specific routine designator> identifies that instance SQL-invoked method.

- 4) If <routine type> specifies ROUTINE, then there shall be exactly one SQL-invoked routine in the schema identified by *SCN* whose <schema qualified routine name> is *RN*. The <specific routine designator> identifies that SQL-invoked routine.

- 4) If FUNCTION is specified, then the SQL-invoked routine that is identified shall be an SQL-invoked regular function. If PROCEDURE is specified, then the SQL-invoked routine that is identified shall be an SQL-invoked procedure. If STATIC METHOD is specified, then the SQL-invoked routine that is identified shall be a static SQL-invoked method. If CONSTRUCTOR METHOD is specified, then the SQL-invoked routine shall be an SQL-invoked constructor method. If INSTANCE METHOD is specified or implicit, then the SQL-invoked routine shall be an instance SQL-invoked method. If ROUTINE is specified, then the SQL-invoked routine that is identified is either an SQL-invoked function or an SQL-invoked procedure.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <specific routine designator> that contains a <routine type> that immediately contains METHOD.

10.7 <collate clause>

Function

Specify a default collation.

Format

```
<collate clause> ::=  
    COLLATE <collation name>
```

Syntax Rules

- 1) Let *C* be the <collation name> contained in the <collate clause>. The schema identified by the explicit or implicit qualifier of the <collation name> shall include the descriptor of *C*.

Access Rules

- 1) Case:
 - a) If <collate clause> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include USAGE on *C*.
 - b) Otherwise, the current privileges shall include USAGE on *C*.

General Rules

None.

Conformance Rules

- 1) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <collate clause>.

10.8 <constraint name definition> and <constraint characteristics>

Function

Specify the name of a constraint and its characteristics.

Format

```
<constraint name definition> ::=
    CONSTRAINT <constraint name>

<constraint characteristics> ::=
    <constraint check time> [ [ NOT ] DEFERRABLE ]
    | [ NOT ] DEFERRABLE [ <constraint check time> ]

<constraint check time> ::=
    INITIALLY DEFERRED
    | INITIALLY IMMEDIATE
```

Syntax Rules

- 1) If <constraint check time> is not specified, then INITIALLY IMMEDIATE is implicit.
- 2) Case:
 - a) If INITIALLY DEFERRED is specified, then:
 - i) NOT DEFERRABLE shall not be specified.
 - ii) If DEFERRABLE is not specified, then DEFERRABLE is implicit.
 - b) If INITIALLY IMMEDIATE is specified or implicit and neither DEFERRABLE nor NOT DEFERRABLE is specified, then NOT DEFERRABLE is implicit.

Access Rules

None.

General Rules

- 1) Let *C* be the constraint identified by <constraint name>.
- 2) If NOT DEFERRABLE is specified, then *C* is not deferrable; otherwise it is deferrable.
- 3) If <constraint check time> is INITIALLY DEFERRED, then the initial constraint mode for *C* is *deferred*; otherwise, the initial constraint mode for *C* is *immediate*.

Conformance Rules

- 1) Without Feature F721, “Deferrable constraints”, conforming SQL language shall not contain a <constraint characteristics>.

NOTE 257 — This means that INITIALLY IMMEDIATE NOT DEFERRABLE is implicit.

- 2) Without Feature F491, “Constraint management”, conforming SQL language shall not contain a <constraint name definition>.

10.9 <aggregate function>

This Subclause is modified by Subclause 11.2, “<aggregate function>”, in ISO/IEC 9075-14.

Function

Specify a value computed from a collection of rows.

Format

```
14 <aggregate function> ::=
    COUNT <left paren> <asterisk> <right paren> [ <filter clause> ]
  | <general set function> [ <filter clause> ]
  | <binary set function> [ <filter clause> ]
  | <ordered set function> [ <filter clause> ]
  | <array aggregate function> [ <filter clause> ]

<general set function> ::=
    <set function type> <left paren> [ <set quantifier> ]
      <value expression> <right paren>

<set function type> ::=
    <computational operation>

<computational operation> ::=
    AVG
  | MAX
  | MIN
  | SUM
  | EVERY
  | ANY
  | SOME
  | COUNT
  | STDDEV_POP
  | STDDEV_SAMP
  | VAR_SAMP
  | VAR_POP
  | COLLECT
  | FUSION
  | INTERSECTION

<set quantifier> ::=
    DISTINCT
  | ALL

<filter clause> ::=
    FILTER <left paren> WHERE <search condition> <right paren>

<binary set function> ::=
    <binary set function type> <left paren> <dependent variable expression> <comma>
      <independent variable expression> <right paren>

<binary set function type> ::=
    COVAR_POP
  | COVAR_SAMP
```

ISO/IEC 9075-2:2008(E)
10.9 <aggregate function>

```
| CORR  
| REGR_SLOPE  
| REGR_INTERCEPT  
| REGR_COUNT  
| REGR_R2  
| REGR_AVGX  
| REGR_AVGY  
| REGR_SXX  
| REGR_SYY  
| REGR_SXY
```

<dependent variable expression> ::=
 <numeric value expression>

<independent variable expression> ::=
 <numeric value expression>

<ordered set function> ::=
 <hypothetical set function>
 | <inverse distribution function>

<hypothetical set function> ::=
 <rank function type> <left paren>
 <hypothetical set function value expression list> <right paren>
 <within group specification>

<within group specification> ::=
 WITHIN GROUP <left paren> ORDER BY <sort specification list> <right paren>

<hypothetical set function value expression list> ::=
 <value expression> [{ <comma> <value expression> }...]

<inverse distribution function> ::=
 <inverse distribution function type> <left paren>
 <inverse distribution function argument> <right paren>
 <within group specification>

<inverse distribution function argument> ::=
 <numeric value expression>

<inverse distribution function type> ::=
 PERCENTILE_CONT
 | PERCENTILE_DISC

<array aggregate function> ::=
 ARRAY_AGG
 <left paren> <value expression> [ORDER BY <sort specification list>] <right paren>

Syntax Rules

- 1) Let *AF* be the <aggregate function>.
- 2) If STDDEV_POP, STDDEV_SAMP, VAR_POP, or VAR_SAMP is specified, then <set quantifier> shall not be specified.
- 3) If <general set function> is specified and <set quantifier> is not specified, then ALL is implicit.
- 4) The argument source of an <aggregate function> is

Case:

- a) If *AF* is immediately contained in a <set function specification> *SFS*, then a group of a grouped table of the aggregation query of *SFS*.
 - b) Otherwise, the collection of rows in the current row's window frame defined by the window structure descriptor identified by the <window function> that simply contains *AF*, as defined in Subclause 7.11, "<window clause>".
- 5) Let *T* be the argument source of *AF*.
- 6) If COUNT is specified, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
- 7) If <general set function> is specified, then:
- a) The <value expression> *VE* shall not contain a <window function>.
 - b) Let *DT* be the declared type of the <value expression>.
 - c) If *AF* specifies a <general set function> whose <set quantifier> is DISTINCT, then *VE* is an operand of a grouping operation. The Syntax Rules of Subclause 9.10, "Grouping operations", apply.
 - d) If *AF* specifies a <set function type> that is MAX or MIN, then *VE* is an operand of an ordering operation. The Syntax Rules of Subclause 9.12, "Ordering operations", apply.
 - e) If EVERY, ANY, or SOME is specified, then *DT* shall be boolean and the declared type of the result is boolean.
 - f) If MAX or MIN is specified, then the declared type of the result is *DT*.
 - g) If SUM or AVG is specified, then:
 - i) *DT* shall be a numeric type or an interval type.
 - ii) If SUM is specified and *DT* is exact numeric with scale *S*, then the declared type of the result is an implementation-defined exact numeric type with scale *S*.
 - iii) If AVG is specified and *DT* is exact numeric, then the declared type of the result is an implementation-defined exact numeric type with precision not less than the precision of *DT* and scale not less than the scale of *DT*.
 - iv) If *DT* is approximate numeric, then the declared type of the result is an implementation-defined approximate numeric with precision not less than the precision of *DT*.
 - v) If *DT* is interval, then the declared type of the result is interval with the same precision as *DT*.
 - h) If VAR_POP or VAR_SAMP is specified, then the declared type of the result is an implementation-defined approximate numeric type. If *DT* is an approximate numeric type, then the precision of the result is not less than the precision of *DT*.
 - i) $\text{STDDEV_POP}(X)$ is equivalent to $\text{SQRT}(\text{VAR_POP}(X))$.
 - j) $\text{STDDEV_SAMP}(X)$ is equivalent to $\text{SQRT}(\text{VAR_SAMP}(X))$.
 - k) If COLLECT is specified, then the declared type of the result is *DT* MULTISSET.
 - l) $\text{COLLECT}(X)$ is equivalent to $\text{FUSION}(\text{MULTISSET}[X])$.

- m) If FUSION is specified, then *DT* shall be a multiset type, and DISTINCT shall not be specified. The declared type of the result is *DT*.
 - n) If INTERSECTION is specified, then *DT* shall be a multiset type, and DISTINCT shall not be specified. *VE* is a multiset operand of a multiset element grouping operation, and the Syntax Rules of Subclause 9.11, “Multiset element grouping operations”, apply. The declared type of the result is *DT*.
- 8) A <filter clause> shall not contain a <query expression>, a <window function>, or an outer reference.
- 9) If <binary set function> is specified, then:
- a) The <dependent variable expression> *DVE* and the <independent variable expression> *IVE* shall not contain a <window function>.
 - b) Let *DTDVE* be the declared type of *DVE* and let *DTIVE* be the declared type of *IVE*.
 - c) Case:
 - i) If REGR_COUNT is specified, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
 - ii) Otherwise, the declared type of the result is an implementation-defined approximate numeric type. If *DTDVE* is an approximate numeric type, then the precision of the result is not less than the precision of *DTDVE*. If *DTIVE* is an approximate numeric type, then the precision of the result is not less than the precision of *DTIVE*.
- 10) If <hypothetical set function> is specified, then:
- a) The <hypothetical set function> shall not contain a <window function>, a <set function specification>, or a <query expression>.
 - b) The number of <value expression>s simply contained in <hypothetical set function value expression list> shall be the same as the number of <sort key>s simply contained in the <sort specification list>.
 - c) For each <value expression> *HSFVE* simply contained in the <hypothetical set function value expression list>, let *SK* be the corresponding <sort key> simply contained in the <sort specification list>.
- Case:
- i) If the declared type of *HSFVE* is a character string type, then the declared type of *SK* shall be a character string type with the same character repertoire as that of *HSFVE*. Syntax Rules of Subclause 9.13, “Collation determination”, are applied with *HSFVE* and *SK* as *VALSET*, yielding the collation returned from *COLL*.
 - ii) Otherwise the declared types of *HSFVE* and *SK* shall be compatible.
- d) Case:
- i) If RANK or DENSE_RANK is specified, then the declared type of the result is exact numeric with implementation-defined precision and with scale 0 (zero).
 - ii) Otherwise, the declared type of the result is approximate numeric with implementation-defined precision.
- 11) If <inverse distribution function> is specified, then:
- a) The <within group specification> shall contain a single <sort specification>.

- b) The <inverse distribution function> shall not contain a <window function>, a <set function specification>, or a <query expression>.
- c) Let *DT* be the declared type of the <value expression> simply contained in the <sort specification>.
- d) *DT* shall be numeric or interval.
- e) The declared type of the result is

Case:

- i) If *DT* is numeric, then approximate numeric with implementation-defined precision.
- ii) If *DT* is interval, then *DT*.

12) If <array aggregate function> *AAG* is specified, then:

- a) Let *VET* be the declared type of the <value expression> immediately contained in *AAG*.
- b) The declared type of *AAG* is array with element type *VET* and maximum cardinality equal to the implementation-defined maximum cardinality *IDMC* for such array types.

Access Rules

None.

General Rules

- 1) If, during the computation of the result of *AF*, an intermediate result is not representable in the declared type of the site that contains that intermediate result, then

Case:

- a) If the most specific type of the result of *AF* is an interval type, then an exception condition is raised: *data exception — interval value out of range*.
- b) 14 If the most specific type of the result of *AF* is a multiset type, then an exception condition is raised: *data exception — multiset value overflow*.
- c) If the most specific type of the result of *AF* is an array type, then an exception condition is raised: *data exception — array data, right truncation*.
- d) Otherwise, an exception condition is raised: *data exception — numeric value out of range*.

2) Case:

- a) If <filter clause> is specified, then the <search condition> is effectively evaluated for each row of *T*. Let *T1* be the collection of rows of *T* for which the result of the <search condition> is True.
- b) Otherwise, let *T1* be *T*.

3) If COUNT(*) is specified, then the result is the cardinality of *T1*.

4) If <general set function> is specified, then:

10.9 <aggregate function>

- a) Let *TX* be the single-column table that is the result of applying the <value expression> to each row of *T1* and eliminating null values. If one or more null values are eliminated, then a completion condition is raised: *warning — null value eliminated in set function*.
- b) Case:
 - i) If DISTINCT is specified, then let *TXA* be the result of eliminating redundant duplicate values from *TX*, using the comparison rules specified in Subclause 8.2, “<comparison predicate>”, to identify the redundant duplicate values.
 - ii) Otherwise, let *TXA* be *TX*.
- c) Let *N* be the cardinality of *TXA*.
- d) Case:
 - i) If COUNT is specified, then the result is *N*.
 - ii) If *TXA* is empty, then the result is the null value.
 - iii) If AVG is specified, then the result is the average of the values in *TXA*.
 - iv) If MAX or MIN is specified, then the result is respectively the maximum or minimum value in *TXA*. These results are determined using the comparison rules specified in Subclause 8.2, “<comparison predicate>”. If *DT* is a user-defined type and the comparison of two values in *TXA* results in *Unknown*, then the maximum or minimum of *TXA* is implementation-dependent.
 - v) If SUM is specified, then the result is the sum of the values in *TXA*. If the sum is not within the range of the declared type of the result, then an exception condition is raised: *data exception — numeric value out of range*.
 - vi) If EVERY is specified, then
 - Case:
 - 1) If the value of some element of *TXA* is *False*, then the result is *False*.
 - 2) Otherwise, the result is *True*.
 - vii) If ANY or SOME is specified, then
 - Case:
 - 1) If the value of some element of *TXA* is *True*, then the result is *True*.
 - 2) Otherwise, the result is *False*.
 - viii) If VAR_POP or VAR_SAMP is specified, then let *S1* be the sum of values in the column of *TXA*, and *S2* be the sum of the squares of the values in the column of *TXA*.
 - 1) If VAR_POP is specified, then the result is $(S2 - S1 * S1 / N) / N$.
 - 2) If VAR_SAMP is specified, then
 - Case:
 - A) If *N* is 1 (one), then the result is the null value.
 - B) Otherwise, the result is $(S2 - S1 * S1 / N) / (N - 1)$

- ix) If FUSION is specified, then the result is the multiset M such that for each value V in the element type of DT , including the null value, the number of elements of M that are identical to V is the sum of the number of identical copies of V in the multisets that are the values of the column in each row of TXA .
- x) If INTERSECTION is specified, then the result is a multiset M such that for each value V in the element type of DT , including the null value, the number of duplicates of V in M is the minimum of the number of duplicates of V in the multisets that are the values of the column in each row of TXA .

NOTE 258 — This rule says “the result is a multiset” rather than “the result is the multiset” because the precise duplicate values are not specified. Thus this calculation is non-deterministic for certain element types, namely those based on character string, datetime with time zone and user-defined types.

NOTE 259 — The notion of one data type being based on another data type is defined in Subclause 4.1, “Data types”.

5) If <binary set function type> is specified, then:

- a) Let TXA be the two-column table that is the result of applying the <dependent variable expression> and the <independent variable expression> to each row of TI and eliminating each row in which either <dependent variable expression> or <independent variable expression> is the null value. If one or more null values are eliminated, then a completion condition is raised: *warning — null value eliminated in set function*.
- b) Let N be the cardinality of TXA , let $SUMX$ be the sum of the column of values of <independent variable expression>, let $SUMY$ be the sum of the column of values of <dependent variable expression>, let $SUMX2$ be the sum of the squares of values in the <independent variable expression> column, let $SUMY2$ be the sum of the squares of values in the <dependent variable expression> column, and let $SUMXY$ be the sum of the row-wise products of the value in the <independent variable expression> column times the value in the <dependent variable expression> column.
- c) Case:
 - i) If REGR_COUNT is specified, then the result is N .
 - ii) If N is 0 (zero), then the result is the null value.
 - iii) If REGR_SXX is specified, then the result is $(SUMX2 - SUMX * SUMX / N)$.
 - iv) If REGR_SYY is specified, then the result is $(SUMY2 - SUMY * SUMY / N)$.
 - v) If REGR_SXY is specified, then the result is $(SUMXY - SUMX * SUMY / N)$.
 - vi) If REGR_AVGX is specified, then the result is $SUMX / N$.
 - vii) If REGR_AVGY is specified, then the result is $SUMY / N$.
 - viii) If COVAR_POP is specified, then the result is $(SUMXY - SUMX * SUMY / N) / N$.
 - ix) If COVAR_SAMP is specified, then
 - Case:
 - 1) If N is 1 (one), then the result is the null value.
 - 2) Otherwise, the result is $(SUMXY - SUMX * SUMY / N) / (N - 1)$
 - x) If CORR is specified, then

Case:

- 1) If $N * SUMX2$ equals $SUMX * SUMX$, then the result is the null value.

NOTE 260 — In this case, all remaining values of <independent variable expression> are equal, and consequently the <independent variable expression> does not correlate with the <dependent variable expression>.

- 2) If $N * SUMY2$ equals $SUMY * SUMY$, then the result is the null value.

NOTE 261 — In this case, all remaining values of <dependent variable expression> are equal, and consequently the <dependent variable expression> does not correlate with the <independent variable expression>.

- 3) Otherwise, the result is $SQRT (POWER (N * SUMXY - SUMX * SUMY , 2) / ((N * SUMX2 - SUMX * SUMX) * (N * SUMY2 - SUMY * SUMY)))$. If the exponent of the approximate mathematical result of the operation is not within the implementation-defined exponent range for the result data type, then the result is the null value.

- xi) If REGR_R2 is specified, then

Case:

- 1) If $N * SUMX2$ equals $SUMX * SUMX$, then the result is the null value.

NOTE 262 — In this case, all remaining values of <independent variable expression> are equal, and consequently the least-squares fit line would be vertical, or, if $N = 1$ (one), there is no uniquely determined least-squares-fit line.

- 2) If $N * SUMY2$ equals $SUMY * SUMY$, then the result is 1 (one).

NOTE 263 — In this case, all remaining values of <dependent variable expression> are equal, and consequently the least-squares fit line is horizontal.

- 3) Otherwise, the result is $POWER (N * SUMXY - SUMX * SUMY , 2) / ((N * SUMX2 - SUMX * SUMX) * (N * SUMY2 - SUMY * SUMY))$. If the exponent of the approximate mathematical result of the operation is not within the implementation-defined exponent range for the result data type, then the result is the null value.

- xii) If REGR_SLOPE(Y, X) is specified, then

Case:

- 1) If $N * SUMX2$ equals $SUMX * SUMX$, then the result is the null value.

NOTE 264 — In this case, all remaining values of <independent variable expression> are equal, and consequently the least-squares fit line would be vertical, or, if $N = 1$ (one), then there is no uniquely determined least-squares-fit line.

- 2) Otherwise, the result is $(N * SUMXY - SUMX * SUMY) / (N * SUMX2 - SUMX * SUMX)$. If the exponent of the approximate mathematical result of the operation is not within the implementation-defined exponent range for the result data type, then the result is the null value.

- xiii) If REGR_INTERCEPT is specified, then

Case:

- 1) If $N * SUMX2$ equals $SUMX * SUMX$, then the result is the null value.

NOTE 265 — In this case, all remaining values of <independent variable expression> are equal, and consequently the least-squares fit line would be vertical, or, if $N = 1$ (one), then there is no uniquely determined least-squares-fit line.

- 2) Otherwise, the result is $(SUMY * SUMX2 - SUMX * SUMXY) / (N * SUMX2 - SUMX * SUMX)$. If the exponent of the approximate mathematical result of the operation is not within the implementation-defined exponent range for the result data type, then the result is the null value.

6) If <hypothetical set function> is specified, then

- a) Let *WIFT* be the <rank function type>.
- b) Let *TNAME* be an implementation-dependent name for *TI*.
- c) Let *K* be the number of <value expression>s simply contained in <hypothetical set function value expression list>.
- d) Let *VE*₁, ..., *VE*_{*K*} be the <value expression>s simply contained in the <hypothetical set function value expression list>.
- e) Let *WIFTVAL*, *MARKER* and *CN*₁, ..., *CN*_{*K*} be distinct implementation-dependent column names.
- f) Let *SP*₁, ..., *SP*_{*K*} be the <sort specification>s simply contained in the <sort specification list>. For each *i*, let *WSP*_{*i*} be the <sort specification> obtained from *SP*_{*i*} by replacing the <sort key> with *CN*_{*i*}.
- g) The result is the result of the <scalar subquery>

```
( SELECT WIFTVAL
  FROM ( SELECT MARKER, WIFT() OVER
          ( ORDER BY WSP1, ..., WSPK )
        FROM ( SELECT 0, SK1, ..., SKK
              FROM TNAME
              UNION ALL
              VALUES (1, VE1, ..., VEK) )
          AS TXNAME (MARKER, CN1, ..., CNK )
        ) AS TEMPTABLE (MARKER, WIFTVAL)
 WHERE MARKER = 1 )
```

7) If <inverse distribution function> is specified, then

- a) Let *NVE* be the value of the <inverse distribution function argument>.
- b) If *NVE* is the null value, then the result is the null value.
- c) If *NVE* is less than 0 (zero) or greater than 1 (one), then an exception condition is raised: *data exception — numeric value out of range*.
- d) Let *TXA* be the single-column table that is the result of applying the <value expression> simply contained in the <sort specification> to each row of *TI* and eliminating null values. If one or more null values are eliminated, then a completion condition is raised: *warning — null value eliminated in set function*. *TXA* is ordered by the <sort specification> as specified in the General Rules of [Subclause 10.10](#), “<sort specification list>”.
- e) Let *TXANAME* be an implementation-dependent name for *TXA*.
- f) Let *TXCOLNAME* be an implementation-dependent column name for the column of *TXA*.
- g) Let *WSP* be obtained from the <sort specification> by replacing the <sort key> with *TXCOLNAME*.

h) Case:

i) If PERCENTILE_CONT is specified, then:

- 1) Let *ROW0* be the greatest exact numeric value with scale 0 (zero) that is less than or equal to $NVE*(N-1)$. Let *ROWLIT0* be a <literal> representing *ROW0*.
- 2) Let *ROW1* be the least exact numeric value with scale 0 (zero) that is greater than or equal to $NVE*(N-1)$. Let *ROWLIT1* be a <literal> representing *ROW1*.
- 3) Let *FACTOR* be an <approximate numeric literal> representing $NVE*(N-1)-ROW0$.
- 4) The result is the result of the <scalar subquery>

```
( WITH TEMPTABLE(X, Y) AS
  ( SELECT ROW_NUMBER()
      OVER (ORDER BY WSP) - 1,
      TXCOLNAME
    FROM TXANAME )
SELECT CAST ( T0.Y + FACTOR * (T1.Y - T0.Y) AS DT )
FROM TEMPTABLE T0, TEMPTABLE T1
WHERE T0.ROWNUMBER = ROWLIT0
AND T1.ROWNUMBER = ROWLIT1 )
```

NOTE 266 — Although ROW_NUMBER is nondeterministic, the values of T0.Y and T1.Y are determined by this expression. Note that the only column of *TXA* is completely ordered by *WSP*. If $NVE*(N-1)$ is a whole number, then the rows selected from T0 and T1 are the same and the result is just T0.Y. Otherwise, the subquery performs a linear interpolation between the two consecutive values whose row numbers in the ordered set, seen as proportions of the whole, bound the argument of the PERCENTILE_CONT operator.

ii) If PERCENTILE_DISC is specified, then

- 1) If the <ordering specification> simply contained in *WSP* is DESC, then let *MAXORMIN* be MAX; otherwise let *MAXORMIN* be MIN.
- 2) Let *NVELIT* be a <literal> representing the value of *NVE*.
- 3) The result is the result of the <scalar subquery>

```
( SELECT MAXORMIN (TXCOLNAME)
  FROM ( SELECT TXCOLNAME,
      CUME_DIST() OVER (ORDER BY WSP)
    FROM TXANAME ) AS TEMPTABLE (TXCOLNAME, CUMEDIST)
WHERE CUMEDIST >= NVELIT )
```

8) If <array aggregate function> is specified, then:

- a) If <sort specification list> is specified, then let *K* be the number of <sort key>s; otherwise, let *K* be 0 (zero).
- b) Let *TXA* be the table of *K*+1 columns obtained by applying the <value expression> immediately contained in the <array aggregate function> to each row of *T1* to obtain the first column of *TXA*, and, for all *i*, 1 (one) ≤ *i* ≤ *K*, applying the <value expression> simply contained in the *i*-th <sort key> to each row of *T1* to obtain the (*i*+1)-th column of *TXA*.
- c) Let *TXA* be ordered according to the values of the <sort key>s found in the second through (*K*+1)-th columns of *TXA*. If *K* is 0 (zero), then the ordering of *TXA* is implementation-dependent.
- d) Let *N* be the number of rows in *TXA*.

- e) If N is greater than $IDMC$, then an exception condition is raised: *data exception — array data, right truncation*.
- f) Let R_i , $1 \text{ (one)} \leq i \leq N$, be the rows of TXA according to the ordering of TXA .
- g) Case:
 - i) If TXA is empty, then the result of <array aggregate function> is the null value.
 - ii) Otherwise, the result of <array aggregate function> is an array of N elements such that for all i , $1 \text{ (one)} \leq i \leq N$, the value of the i -th element is the value of the first column of R_i .

NOTE 267 — Null values are not eliminated when computing <array aggregate function>. This, plus the optional <sort specification list>, sets <array aggregate function> apart from <general set function>s.

Conformance Rules

- 1) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <computational operation> that immediately contains EVERY, ANY, or SOME.
- 2) Without Feature F561, “Full value expressions”, or Feature F801, “Full set function”, conforming SQL language shall not contain a <general set function> that immediately contains DISTINCT and contains a <value expression> that is not a column reference.
- 3) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <general set function> that contains a <computational operation> that immediately contains COUNT and does not contain a <set quantifier> that immediately contains DISTINCT.
- 4) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <general set function> that does not contain a <set quantifier> that immediately contains DISTINCT and that contains a <value expression> that contains a column reference that does not reference a column of T .
- 5) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <binary set function> that does not contain either a <dependent variable expression> or an <independent variable expression> that contains a column reference that references a column of T .
- 6) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <value expression> simply contained in a <general set function> that contains a column reference that is an outer reference where the <value expression> is not a column reference.
- 7) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <numeric value expression> simply contained in a <dependent variable expression> or an <independent variable expression> that contains a column reference that is an outer reference and in which the <numeric value expression> is not a column reference.
- 8) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a column reference contained in an <aggregate function> that contains a reference to a column derived from a <value expression> that generally contains an <aggregate function> $SFS2$ without an intervening <routine invocation>.
- 9) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <computational operation> that immediately contains STDDEV_POP, STDDEV_SAMP, VAR_POP, or VAR_SAMP.

10.9 <aggregate function>

- 10) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <binary set function type>.
- 11) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <hypothetical set function> or an <inverse distribution function>.
- 12) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <filter clause>.
- 13) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <computational operation> that immediately contains COLLECT.
- 14) Without Feature S275, “Advanced multiset support”, conforming SQL language shall not contain a <computational operation> that immediately contains FUSION or INTERSECTION.

NOTE 268 — If INTERSECTION is specified, then the Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

- 15) Without Feature T052, “MAX and MIN for row types”, conforming SQL language shall not contain a <computational operation> that immediately contains MAX or MIN in which the declared type of the <value expression> is a row type.

NOTE 269 — If DISTINCT is specified, then the Conformance Rules of Subclause 9.10, “Grouping operations”, also apply. If MAX or MIN is specified, then the Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.

- 16) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain a <hypothetical set function value expression list> or a <sort specification list> that simply contains a <value expression> that contains more than one column reference, one of which is an outer reference.
- 17) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain an <inverse distribution function> that contains an <inverse distribution function argument> or a <sort specification> that contains more than one column reference, one of which is an outer reference.
- 18) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain an <aggregate function> that contains a <general set function> whose simply contained <value expression> contains more than one column reference, one of which is an outer reference.
- 19) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain an <aggregate function> that contains a <binary set function> whose simply contained <dependent variable expression> or <independent variable expression> contains more than one column reference, one of which is an outer reference.
- 20) Without Feature S098, “ARRAY_AGG”, conforming SQL language shall not contain an <array aggregate function>.

10.10 <sort specification list>

Function

Specify a sort order.

Format

```
<sort specification list> ::=
    <sort specification> [ { <comma> <sort specification> }... ]

<sort specification> ::=
    <sort key> [ <ordering specification> ] [ <null ordering> ]

<sort key> ::=
    <value expression>

<ordering specification> ::=
    ASC
    | DESC

<null ordering> ::=
    NULLS FIRST
    | NULLS LAST
```

Syntax Rules

- 1) Let *DT* be the declared type of the <value expression> simply contained in the <sort key> contained in a <sort specification>.
- 2) Each <value expression> simply contained in the <sort key> contained in a <sort specification> is an operand of an ordering operation. The Syntax Rules of [Subclause 9.12, “Ordering operations”](#), apply.
- 3) If <null ordering> is not specified, then an implementation-defined <null ordering> is implicit. The implementation-defined default for <null ordering> shall not depend on the context outside of <sort specification list>.

Access Rules

None.

General Rules

- 1) A <sort specification list> defines an ordering of rows, as follows:
 - a) Let *N* be the number of <sort specification>s.
 - b) Let K_i , $1 \text{ (one)} \leq i \leq N$, be the <sort key> contained in the *i*-th <sort specification>.
 - c) Each <sort specification> specifies the *sort direction* for the corresponding sort key K_i . If DESC is not specified in the *i*-th <sort specification>, then the sort direction for K_i is ascending and the applicable

<comp op> is the <less than operator>. Otherwise, the sort direction for K_i is descending and the applicable <comp op> is the <greater than operator>.

- d) Let P be any row of the collection of rows to be ordered, and let Q be any other row of the same collection of rows.
- e) Let PV_i and QV_i be the values of K_i in P and Q , respectively. The relative position of rows P and Q in the result is determined by comparing PV_i and QV_i according to the rules of [Subclause 8.2](#), “<comparison predicate>” where the <comp op> is the applicable <comp op> for K_i , with the following special treatment of null values.

Case:

- i) If PV_i and QV_i are both the null value, then they are considered equal to each other.
- ii) If PV_i is the null value and QV_i is not the null value, then

Case:

- 1) If NULLS FIRST is specified or implied, then PV_i <comp op> QV_i is considered to be True.
- 2) If NULLS LAST is specified or implied, then PV_i <comp op> QV_i is considered to be False.
- iii) If PV_i is not the null value and QV_i is the null value, then

Case:

- 1) If NULLS FIRST is specified or implied, then PV_i <comp op> QV_i is considered to be False.
- 2) If NULLS LAST is specified or implied, then PV_i <comp op> QV_i is considered to be True.
- f) PV_i is said to *precede* QV_i if the value of the <comparison predicate> “ PV_i <comp op> QV_i ” is True for the applicable <comp op>.
- g) If PV_i and QV_i are not the null value and the result of “ PV_i <comp op> QV_i ” is Unknown, then the relative ordering of PV_i and QV_i is implementation-dependent.
- h) The relative position of row P is before row Q if PV_n precedes QV_n for some n , $1 \text{ (one)} \leq n \leq N$, and PV_i is not distinct from QV_i for all $i < n$.
- i) Two rows that are not distinct with respect to the <sort specification>s are said to be *peers* of each other. The relative ordering of peers is implementation-dependent.

Conformance Rules

- 1) Without Feature T611, “Elementary OLAP operations”, conforming SQL language shall not contain a <null ordering>.

NOTE 270 — The Conformance Rules of [Subclause 9.12](#), “Ordering operations”, also apply.

10.11 Determination of view privileges

Subclause Signature

"Determination of view privileges" [General Rules] (
Parameter: "VIEW"
) Returns: "DESCRIPTOR SET"

Function

Determine the privilege descriptors of the view whose <action> is INSERT, UPDATE, or DELETE. In addition, determine the view privilege dependency descriptors of the view.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *V* be the *VIEW* in an application of this Subclause. The *SVPD* that is the result of the application of this Subclause is *DESCRIPTOR SET*.
- 2) Let *A* be the <authorization identifier> that owns the schema identified by the <table name> of *V*.
- 3) The General Rules of Subclause 10.12, "Determination of view component privileges", are applied with *V* as *VIEW*, yielding a set of view component privilege descriptors returned from *DESCRIPTOR SET*.
- 4) A view component privilege descriptor *SPD* is *simply dependent* on another privilege descriptor *PD* if *SPD* is immediately dependent on *PD*, or if there is a view component privilege descriptor *SPD2* such that *SPD* is immediately dependent on *SPD2* and *SPD2* is simply dependent on *PD*.
- 5) Let *VC_N* be the view component that is the <query expression> immediately contained in the <view definition> of *V*.
 - a) For each view component table privilege descriptor *VCTPD* whose identified object is *VC_N*:
 - i) A privilege descriptor *PD1* is created, as follows: the identified object is *V*, the action is the same as the action of *VCTPD*, the grantor is the special grantor value "_SYSTEM", the grantee is *A*, and the privilege is grantable if and only if *VCTPD* is grantable.
 - ii) For each privilege descriptor *PD* such that *VCTPD* is simply dependent on *PD*, and such that the object of *PD* is not a view component or a column of a view component, a view privilege dependency descriptor is created, as follows: the supporting privilege descriptor is *PD* and the dependent privilege descriptor is *PD1*.

10.11 Determination of view privileges

- b) For each view component column privilege descriptor *VCCPD* whose identified object is VC_N :
 - i) A privilege descriptor *PD2* is created, as follows: the identified object is the column of *V* that is the counterpart of the column identified by *VCCPD*, the action is the same as the action of *VCCPD*, the grantor is the special grantor value “_SYSTEM”, the grantee is *A*, and the privilege is grantable if and only if *VCCPD* is grantable.
 - ii) For each privilege descriptor *PD* such that *VCCPD* is simply dependent on *PD*, and such that the object of *PD* is not a view component or a column of a view component, a view privilege dependency descriptor is created, as follows: the supporting privilege descriptor is *PD* and the dependent privilege descriptor is *PD2*.
- 6) If, for every column *C* of *V*, there is a column privilege descriptor *CPD* whose identified object is *C*, action is UPDATE, grantor is the special grantor value “_SYSTEM”, and grantee is *A*, then it is implementation-defined whether a table privilege descriptor *TPD* is created whose identified object is *V*, action is UPDATE, grantor is the special grantor value “_SYSTEM”, grantee is *A*. If such a table privilege descriptor is created, then it is directly dependent on every such column privilege descriptor, and it has an indication that the privilege is grantable if every such column privilege descriptor has an indication that it is grantable.
 If *TPD* is created, then a collection of view privilege dependency descriptors is created, one for each column *C* of *V*, in which the supporting privilege descriptor is *CPD* and the dependent privilege descriptor is *TPD*.
- 7) If, for every column *C* of *V*, there is a column privilege descriptor *CPD* whose identified object is *C*, action is INSERT, grantor is the special grantor value “_SYSTEM”, grantee is *A*, then it is implementation-defined whether a table privilege descriptor *TPD* is created whose identified object is *V*, action is INSERT, grantor is the special grantor value “_SYSTEM”, grantee is *A*. If such a table privilege descriptor is created, then it is directly dependent on every such column privilege descriptor, and it has an indication that the privilege is grantable if every such column privilege descriptor has an indication that it is grantable.
 If *TPD* is created, then a collection of view privilege dependency descriptors is created, one for each column *C* of *V*, in which the supporting privilege descriptor is *CPD* and the dependent privilege descriptor is *TPD*.
- 8) All view component privilege descriptors are destroyed.
- 9) The set *SVPD* of view privilege descriptors and view privilege dependency descriptors created by the General Rules of this Subclause is the result of the application of this Subclause.

Conformance Rules

None.

10.12 Determination of view component privileges

Subclause Signature

"Determination of view component privileges" [General Rules] (
 Parameter: "VIEW"
) Returns: "DESCRIPTOR SET"

Function

Determine view component privilege descriptors for all view components of a view.

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *V* be the *VIEW* in an application of this Subclause. The *SVCPD* that is the result of the application of this Subclause is *DESCRIPTOR SET*.
- 2) Let *A* be the <authorization identifier> that owns the schema identified by the <table name> of *V*.
- 3) Let *VC*₁, ..., *VC*_{*N*} be an enumeration of the view components of *V*. The enumeration shall have the following properties:
 - a) For all *i* and *j* between 1 (one) and *N*, if *VC*_{*i*} is included in *VC*_{*j*}, then *i* < *j*.
 - b) For all *i* and *j* between 1 (one) and *N*, if *VC*_{*i*} is a <query expression> simply contained in a <with list element> and *VC*_{*j*} references the table defined by *VC*_{*i*}, then *i* < *j*.

NOTE 271 — A depth-first left-to-right traversal of the BNF of the <view definition> of *V* is one way to obtain such an enumeration.

- 4) If *V* effectively updatable, then the following subrules are performed to recursively create certain view component privilege descriptors. The following subrules also recursively define when a view component privilege is *immediately dependent* on another privilege descriptor.

For each *i* between 1 (one) and *N*, in that order,

Case:

- a) If *VC*_{*i*} is an updatable <query specification>, then:
 - i) If the <from clause> *VC*_{*i*} contains exactly one <table reference> *TR*, then

- 1) Case:
 - A) If TR is a <table name> or <transition table name>, then let S be the set of applicable privileges for A on the table referenced by TR .
 - B) Otherwise, TR is a <query name> and references some VC_j , where $j < i$. Let S be the set of view component privilege descriptors whose identified object is VC_j or a column of VC_j .
 - 2) If S contains a table privilege descriptor PD whose action is DELETE, then a view component table privilege descriptor is created as follows: the identified object is VC_i , the action is DELETE, the grantor is the special grantor value “_SYSTEM”, and the grantee is A . The privilege is grantable if and only if PD indicates a grantable privilege.
 - 3) For each updatable column C of VC_i , let CC be the counterpart of C in the table identified by TR .
 - A) If S contains a column privilege descriptor PD whose action is UPDATE(CC), then a view component table privilege descriptor is created as follows: the identified object is C , the action is UPDATE, the grantor is the special grantor value “_SYSTEM”, and the grantee is A . The privilege is grantable if and only if PD indicates a grantable privilege. The privilege descriptor is immediately dependent on PD .
 - B) If V is insertable-into, and S contains a column privilege descriptor PD whose action is INSERT(CC), then a view component table privilege descriptor is created as follows: the identified object is C , the action is INSERT, the grantor is the special grantor value “_SYSTEM”, and the grantee is A . The privilege is grantable if and only if PD indicates a grantable privilege. The privilege descriptor is immediately dependent on PD .
- ii) Otherwise:
- 1) If, for every leaf underlying table LUT of VC_i such that VC_i is one-to-one with LUT , the applicable privileges for A include DELETE on LUT , then a view component table privilege descriptor is created whose identified object is VC_i , action is DELETE, grantor is the special grantor value “_SYSTEM”, and the grantee is A . The privilege is grantable if and only if the applicable privileges for A includes grantable DELETE privilege on each such LUT . The privilege descriptor is immediately dependent on every privilege descriptor whose identified object is such a leaf underlying table, the action is DELETE, and grantee is A .
 - 2) For each updatable column C of VC_i , let LUT be the leaf underlying table of VC_i that has a counterpart CC to C .
 - A) If the applicable privileges for A include UPDATE(CC) on LUT , then a view component column privilege descriptor $VCCPD$ is created, as follows: the identified object is C , the action is UPDATE, the grantor is the special grantor value “_SYSTEM”, and the grantee is A . The privilege is grantable if and only if the applicable privilege for A includes grantable UPDATE(CC) privilege on LUT . The privilege descriptor is immediately dependent on every privilege descriptor whose identified object is CC , action is UPDATE, and grantee is A .
 - B) If V is insertable-into, and the applicable privileges for A include INSERT (CC) on LUT , then a view component column privilege descriptor $VCCPD$ is created, as follows: the identified object is C , the action is INSERT, the grantor is the special grantor value

10.12 Determination of view component privileges

“_SYSTEM”, and the grantee is *A*. The privilege is grantable if and only if the applicable privilege for *A* includes grantable INSERT(*CC*) privilege on *LUT*. The privilege descriptor is immediately dependent on every privilege descriptor whose identified object is *CC*, action is INSERT, and grantee is *A*.

- b) If VC_i is a <table value constructor>, then there are no view component privilege descriptors that identify VC_i as object.
- c) If VC_i is a <table value constructor>, then let *T* be the table identified by the <explicit table>.
 - i) If the applicable privileges for *A* include DELETE on *T*, then a view component table privilege descriptor is created whose identified object is VC_i , action is DELETE, grantor is the special grantor value “_SYSTEM”, and the grantee is *A*. The privilege is grantable if and only if the applicable privileges for *A* includes grantable DELETE privilege on *T*. The privilege descriptor is immediately dependent on the privilege descriptor whose identified object is *T*, action is DELETE, and grantee is *A*.
 - ii) For each updatable column *C* of VC_i , let *CC* be the counterpart to *C* in *T*. If the applicable privileges for *A* include UPDATE(*CC*) on *T*, then a view component column privilege descriptor *VCCPD* is created, as follows: the identified object is *C*, the action is UPDATE, the grantor is the special grantor value “_SYSTEM”, and the grantee is *A*. The privilege is grantable if and only if the applicable privilege for *A* includes grantable UPDATE(*CC*) privilege on *T*. The privilege descriptor is immediately dependent on the privilege descriptor whose identified object is *CC*, action is UPDATE, and grantee is *A*.
 - iii) For each updatable column *C* of VC_i , let *CC* be the counterpart to *C* in *T*. If the applicable privileges for *A* include INSERT(*CC*) on *T*, then a view component column privilege descriptor *VCCPD* is created, as follows: the identified object is *C*, the action is INSERT, the grantor is the special grantor value “_SYSTEM”, and the grantee is *A*. The privilege is grantable if and only if the applicable privilege for *A* includes grantable INSERT(*CC*) privilege on *T*. The privilege descriptor is immediately dependent on the privilege descriptor whose identified object is *CC*, action is INSERT, and grantee is *A*.
- d) If VC_i is a <query expression>, then

Case:

- i) If VC_i is a <simple table> *ST*, then there exists a $j < i$ such that *ST* is VC_j .
 - 1) For each view component table privilege descriptor *VCTPD* of VC_j , a new view component table privilege descriptor is created, as follows: the identified object is VC_i , the grantor is the special grantor value “_SYSTEM”, the grantee is *A*, and the action, and the indication of whether the privilege is grantable is the same as in *VCTPD*. The privilege descriptor is immediately dependent on *VCTPD*.
 - 2) For each view component column privilege descriptor *VCCPD* of VC_j , a new view component column privilege descriptor of VC_i is created, as follows: the identified object is the column of VC_i that is the counterpart of the column of VC_j identified by *VCCPD*, the grantor is the special grantor value “_SYSTEM”, the grantee is *A*, and the the action and the indication of whether the privilege is grantable are the same as in *VCCPD*. The privilege descriptor is immediately dependent on *VCCPD*.

10.12 Determination of view component privileges

- ii) Otherwise, VC_i immediately contains UNION ALL. Let VC_l and VC_r be the left and right operands of VC_i , respectively.
 - 1) If there is a view component table privilege descriptor $VCTPD_l$ whose identified object is VC_l and whose action is DELETE, and there is a view component table privilege descriptor $VCTPD_r$ whose identified object is VC_r and whose action is DELETE, then a new view component table privilege descriptor is created as follows: the identified object is VC_i , the action is DELETE, the grantor is the special grantor value “_SYSTEM”, the grantee is A, and the privilege is grantable if and only if both $VCTPD_l$ and $VCTPD_r$ indicate that the privilege is grantable. The privilege descriptor is immediately dependent on $VCTPD_l$ and $VCTPD_r$.
 - 2) For each updatable column C of VC_i , let C_l and C_r be the counterparts of C in VC_l and VC_r , respectively. If there is a view component column privilege descriptor $VCTPD_l$ whose identified object is C_l and whose action is UPDATE, and there is a view component column privilege descriptor $VCTPD_r$ whose identified object is C_r and whose action is UPDATE, then a new view component table privilege descriptor is created as follows: the identified object is C , the action is UPDATE, the grantor is the special grantor value “_SYSTEM”, the grantee is A, and the privilege is grantable if and only if both $VCTPD_l$ and $VCTPD_r$ indicate that the privilege is grantable. The privilege descriptor is immediately dependent on $VCTPD_l$ and $VCTPD_r$.
- 5) The set $SVCPD$ of view component table privilege descriptors and view component column privilege descriptors created by the General Rules of this Subclause is the result of the application of this Subclause.

Conformance Rules

None.

11 Schema definition and manipulation

11.1 <schema definition>

This Subclause is modified by Subclause 9.1, “<schema definition>”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 11.1, “<schema definition>”, in ISO/IEC 9075-9.

Function

Define a schema.

Format

```
<schema definition> ::=
    CREATE SCHEMA <schema name clause>
        [ <schema character set or path> ]
        [ <schema element>... ]

<schema character set or path> ::=
    <schema character set specification>
    | <schema path specification>
    | <schema character set specification> <schema path specification>
    | <schema path specification> <schema character set specification>

<schema name clause> ::=
    <schema name>
    | AUTHORIZATION <schema authorization identifier>
    | <schema name> AUTHORIZATION <schema authorization identifier>

<schema authorization identifier> ::=
    <authorization identifier>

<schema character set specification> ::=
    DEFAULT CHARACTER SET <character set specification>

<schema path specification> ::=
    <path specification>

04 09 <schema element> ::=
    <table definition>
    | <view definition>
    | <domain definition>
    | <character set definition>
    | <collation definition>
    | <transliteration definition>
    | <assertion definition>
    | <trigger definition>
    | <user-defined type definition>
```

```
| <user-defined cast definition>  
| <user-defined ordering definition>  
| <transform definition>  
| <schema routine>  
| <sequence generator definition>  
| <grant statement>  
| <role definition>
```

Syntax Rules

- 1) If <schema name> is not specified, then a <schema name> equal to <schema authorization identifier> is implicit.
- 2) If AUTHORIZATION <schema authorization identifier> is not specified, then
Case:
 - a) If the <schema definition> is contained in an SQL-client module that has a <module authorization identifier> specified, then an <authorization identifier> equal to that <module authorization identifier> is implicit for the <schema definition>.
 - b) Otherwise, an <authorization identifier> equal to the SQL-session user identifier is implicit.
- 3) The <unqualified schema name> of the explicit or implicit <schema name> shall not be equivalent to the <unqualified schema name> of the <schema name> of any other schema in the catalog identified by the <catalog name> of <schema name>.
- 4) If a <schema definition> is contained in an <externally-invoked procedure> in an <SQL-client module definition>, then the effective <schema authorization identifier> and <schema name> during processing of the <schema definition> are, respectively, the <schema authorization identifier> and <schema name> specified or implicit in the <schema definition>.

NOTE 272 — Other SQL-statements executed in <externally-invoked procedure>s in the SQL-client module have the <module authorization identifier> and <schema name> specified or implicit for the SQL-client module.
- 5) If <schema character set specification> is not specified, then a <schema character set specification> that specifies an implementation-defined character set that contains at least every character that is in <SQL language character> is implicit.
- 6) If <schema path specification> is not specified, then a <schema path specification> containing an implementation-defined <schema name list> that contains the <schema name> contained in <schema name clause> is implicit.
- 7) The explicit or implicit <catalog name> of each <schema name> contained in the <schema name list> of the <schema path specification> shall be equivalent to the <catalog name> of the <schema name> contained in the <schema name clause>.
- 8) The <schema name list> of the explicit or implicit <schema path specification> is used as the SQL-path of the schema. The SQL-path is used to effectively qualify unqualified <routine name>s that are immediately contained in <routine invocation>s that are contained in the <schema definition>.

NOTE 273 — <routine name> is defined in [Subclause 5.4, “Names and identifiers”](#).

Access Rules

- 1) The privileges necessary to execute the <schema definition> are implementation-defined.

General Rules

- 1) A <schema definition> creates an SQL-schema *S* in a catalog. *S* includes:
 - a) A schema name that is equivalent to the explicit or implicit <schema name>.
 - b) A schema authorization identifier that is equivalent to the explicit or implicit <authorization identifier>.
 - c) A schema character set name that is equivalent to the explicit or implicit <schema character set specification>.
 - d) A schema SQL-path that is equivalent to the explicit or implicit <schema path specification>.
 - e) The descriptor created by every <schema element> of the <schema definition>.
- 2) The owner of *S* is schema authorization identifier.
- 3) The explicit or implicit <character set specification> is used as the default character set used for all <column definition>s and <domain definition>s that do not specify an explicit character set.

Conformance Rules

- 1) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <schema path specification>.
- 2) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <schema character set specification>.
- 3) Without Feature F171, “Multiple schemas per user”, conforming SQL language shall not contain a <schema name clause> that contains a <schema name>.
- 4) Without Feature T332, “Extended roles”, in conforming SQL language a <schema authorization identifier> shall not be a <role name>.

11.2 <drop schema statement>

This Subclause is modified by Subclause 9.2, “<drop schema statement>”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 11.2, “<drop schema statement>”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 9.1, “<drop schema statement>”, in ISO/IEC 9075-13.

Function

Destroy a schema.

Format

```
<drop schema statement> ::=
    DROP SCHEMA <schema name> <drop behavior>

<drop behavior> ::=
    CASCADE
  | RESTRICT
```

Syntax Rules

- 1) Let *S* be the schema identified by <schema name> *SN*.
- 2) *SN* shall identify a schema in the catalog identified by the explicit or implicit <catalog name>.
- 3) *SN* shall not identify the schema named “INFORMATION_SCHEMA” in any catalog.
- 4) 09 13 If RESTRICT is specified, then *S* shall not contain any persistent base tables, global temporary tables, created local temporary tables, views, domains, assertions, character sets, collations, transliterations, triggers, user-defined types, SQL-invoked routines, sequence generators, or roles, and the <schema name> of *S* shall not be generally contained in the SQL routine body of any routine descriptor.

NOTE 274 — If CASCADE is specified, then such objects will be dropped by the effective execution of the SQL schema manipulation statements specified in the General Rules of this Subclause.

Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by *SN*.

General Rules

- 1) 09 13 Let *T* be the <table name> included in the descriptor of any base table or temporary table included in *S*. The following <drop table statement> is effectively executed:

```
DROP TABLE T CASCADE
```

- 2) Let *V* be the <table name> included in the descriptor of any view included in *S*. The following <drop view statement> is effectively executed:

DROP VIEW *V* CASCADE

- 3) Let *D* be the <domain name> included in the descriptor of any domain included in *S*. The following <drop domain statement> is effectively executed:

DROP DOMAIN *D* CASCADE

- 4) Let *A* be the <constraint name> included in the descriptor of any assertion included in *S*. The following <drop assertion statement> is effectively executed:

DROP ASSERTION *A* CASCADE

- 5) Let *CD* be the <collation name> included in the descriptor of any collation included in *S*. The following <drop collation statement> is effectively executed:

DROP COLLATION *CD* CASCADE

- 6) Let *TD* be the <transliteration name> included in the descriptor of any transliteration included in *S*. The following <drop transliteration statement> is effectively executed:

DROP TRANSLATION *TD*

- 7) Let *RD* be the <character set name> included in the descriptor of any character set included in *S*. The following <drop character set statement> is effectively executed:

DROP CHARACTER SET *RD*

- 8) 04 Let *DT* be the <user-defined type name> included in the descriptor of any user-defined type included in *S*. The following <drop data type statement> is effectively executed:

DROP TYPE *DT* CASCADE

- 9) Let *TT* be the <trigger name> included in the descriptor of any trigger included in *S*. The following <drop trigger statement> is effectively executed:

DROP TRIGGER *TT*

- 10) For every SQL-invoked routine *R* whose descriptor is included in *S*, let *SN* be the <specific name> of *R*. The following <drop routine statement> is effectively executed for every *R*:

DROP SPECIFIC ROUTINE *SN* CASCADE

- 11) 04 Let *R* be any SQL-invoked routine whose routine descriptor includes an SQL routine body that contains the <schema name> of *S*. Let *SN* be the <specific name> of *R*. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE *SN* CASCADE

- 12) 13 Let *SEQN* be the sequence generator name included in the descriptor of any sequence generator included in *S*. The following <drop sequence generator statement> is effectively executed:

DROP SEQUENCE *SEQN* CASCADE

13) *S* is destroyed.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop schema statement>.

11.3 <table definition>

*This Subclause is modified by Subclause 9.3, “<table definition>”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 11.3, “<table definition>”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 9.2, “<table definition>”, in ISO/IEC 9075-13.*

Function

Define a persistent base table, a created local temporary table, or a global temporary table.

Format

```
<table definition> ::=
    CREATE [ <table scope> ] TABLE <table name> <table contents source>
        [ ON COMMIT <table commit action> ROWS ]

<table contents source> ::=
    <table element list>
    | <typed table clause>
    | <as subquery clause>

<table scope> ::=
    <global or local> TEMPORARY

<global or local> ::=
    GLOBAL
    | LOCAL

<table commit action> ::=
    PRESERVE
    | DELETE

<table element list> ::=
    <left paren> <table element> [ { <comma> <table element> }... ] <right paren>

<table element> ::=
    <column definition>
    | <table constraint definition>
    | <like clause>

<typed table clause> ::=
    OF <path-resolved user-defined type name> [ <subtable clause> ]
        [ <typed table element list> ]

<typed table element list> ::=
    <left paren> <typed table element>
        [ { <comma> <typed table element> }... ] <right paren>

<typed table element> ::=
    <column options>
    | <table constraint definition>
    | <self-referencing column specification>

<self-referencing column specification> ::=
    REF IS <self-referencing column name> [ <reference generation> ]
```

```

<reference generation> ::=
    SYSTEM GENERATED
    | USER GENERATED
    | DERIVED

<self-referencing column name> ::=
    <column name>

<column options> ::=
    <column name> WITH OPTIONS <column option list>

09 <column option list> ::=
    [ <scope clause> ] [ <default clause> ] [ <column constraint definition>... ]

<subtable clause> ::=
    UNDER <supertable clause>

<supertable clause> ::=
    <supertable name>

<supertable name> ::=
    <table name>

<like clause> ::=
    LIKE <table name> [ <like options> ]

<like options> ::=
    <like option>...

<like option> ::=
    <identity option>
    | <column default option>
    | <generation option>

<identity option> ::=
    INCLUDING IDENTITY
    | EXCLUDING IDENTITY

<column default option> ::=
    INCLUDING DEFAULTS
    | EXCLUDING DEFAULTS

<generation option> ::=
    INCLUDING GENERATED
    | EXCLUDING GENERATED

<as subquery clause> ::=
    [ <left paren> <column name list> <right paren> ] AS <table subquery>
    <with or without data>

<with or without data> ::=
    WITH NO DATA
    | WITH DATA

```

Syntax Rules

- 1) 04 The <table contents source> shall not contain a <host parameter specification>, an <SQL parameter reference>, a <dynamic parameter specification>, or an <embedded variable specification>.

- 2) Let T be the table defined by the <table definition> TD . Let TN be the <table name> simply contained in TD .
- 3) If a <table definition> is contained in a <schema definition> SD and TN contains a <local or schema qualifier>, then that <local or schema qualifier> shall be equivalent to the implicit or explicit <schema name> of SD .
- 4) TN shall not identify an existing table descriptor.
- 5) If the <table definition> is contained in a <schema definition>, then let A be the explicit or implicit <authorization identifier> of the <schema definition>. Otherwise, let A be the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of TN .
- 6) If <table element list> TEL is specified, then:
 - a) TEL shall contain at least one <column definition> or <like clause>.
 - b) For each <like clause> LC that is directly contained in TEL :
 - i) Let LT be the table identified by the <table name> contained in LC .
 - ii) If LT is a viewed table, then <like options> shall not be specified.
 - iii) Let D be the degree of LT . For i , $1 \text{ (one)} \leq i \leq D$:
 - 1) Let LCD_i be the column descriptor of the i -th column of LT .
 - 2) Let LCN_i be the column name included in LCD_i .
 - 3) Let LDT_i be the data type included in LCD_i .
 - 4) If the nullability characteristic included in LCD_i is known not nullable, then let LNC_i be NOT NULL; otherwise, let LNC_i be the zero-length string.
 - 5) Let CD_i be the <column definition>

$$LCN_i \ LDT_i \ LNC_i$$
 - iv) If <like options> is specified, then:
 - 1) <identity option> shall not be specified more than once, <column default option> shall not be specified more than once, and <generation option> shall not be specified more than once.
 - 2) If <identity option> is not specified, then EXCLUDING IDENTITY is implicit.
 - 3) If <column default option> is not specified, then EXCLUDING DEFAULTS is implicit.
 - 4) If <generation option> is not specified, then EXCLUDING GENERATED is implicit.
 - 5) If INCLUDING IDENTITY is specified and LT includes an identity column, then let ICD be the column descriptor of that column included in the table descriptor of LT . Let SGD be the sequence generator descriptor included in ICD .
 - A) Let SV be the start value included in ICD .
 - B) Let IV be the increment included in SGD .
 - C) Let MAX be the maximum value included in SGD .

- D) Let MIN be the minimum value included in SGD .
- E) Let CYC be the cycle option included in SGD .
- F) Let k be the ordinal position in which the column described by ICD appears in the table identified by LT .
- G) Case:
 - I) If ICD indicates that values are always generated, then let G be GENERATED ALWAYS.
 - II) If ICD indicates that values are generated by default, then let G be GENERATED BY DEFAULT.
- H) The value of CD_k is replaced by:

LCN_k LDT_k
 G AS IDENTITY (START WITH SV , INCREMENT BY IV ,
 MAXVALUE MAX , MINVALUE MIN , CYC) LNC_k

- 6) If INCLUDING GENERATED is specified, then let GCD_j , $1 \text{ (one)} \leq j \leq D$, be the column descriptors included in the descriptor of LT , with j being the ordinal position of the column described by GCD_j . For each GCD_j that indicates that the column it describes is a generated column:
 - A) Let GE_j be the <generation expression> included in GCD_j , where the <table name> contained in any contained <column reference> is replaced by TN .
 - B) The value of CD_j is replaced by

LCN_j LDT_j GENERATED ALWAYS AS GE_j LCN_j

- 7) If INCLUDING DEFAULTS is specified, then let DCD_m , $1 \text{ (one)} \leq m \leq D$, be the column descriptors included in the descriptor of LT , with m being the ordinal position of the column described by DCD_m .

For each DCD_m , if DCD_m includes a <default option> DO_m , then the value of CD_m is replaced by

LCN_m LDT_m DEFAULT DO_m LCN_m

- v) LC is effectively replaced by:

CD_1, \dots, CD_D

NOTE 275 — <column constraint>s, except for NOT NULL, are not included in CD_i ; <column constraint definition>s are effectively transformed to <table constraint definition>s and are thereby also excluded.

- 7) If <as subquery clause> is specified, then:

- a) Let ASQ be the <table subquery> immediately contained in <as subquery clause>. Let QT be the table specified by ASQ .

- b) If any two columns in QT have equivalent <column name>s, or if any column of QT has an implementation-dependent name, then <column name list> shall be specified.
- c) Let D be the degree of QT .
- d) <column name list> shall not contain two or more equivalent <column name>s.
- e) The number of <column name>s in <column name list> shall be D .
- f) For i , $1 \text{ (one)} \leq i \leq D$:
 - i) Case:
 - 1) If <column name list> is specified, then let QCN_i be the i -th <column name> in that <column name list>.
 - 2) Otherwise, let QCN_i be the <column name> of the i -th column of QT .
 - ii) Let QDT_i be the declared type of the i -th column of QT .
 - iii) If the nullability characteristic of the i -th column of QT is known not nullable, then let QNC_i be NOT NULL; otherwise, let QNC_i be the zero-length string.
 - iv) Let CD_i be the <column definition>

$$QCN_i \quad QDT_i \quad QNC_i$$

- g) <as subquery clause> is effectively replaced by a <table element list> TEL of the form:

$$CD_1, \dots, CD_D$$

- 8) If <typed table clause> TTC is specified, then:

- a) The <user-defined type name> simply contained in <path-resolved user-defined type name> shall identify a structured type ST .
- b) If <subtable clause> is specified, then <self-referencing column specification> shall not be specified. Otherwise, <self-referencing column specification> shall be specified exactly once.
- c) If <self-referencing column specification> $SRCS$ is specified, then let RST be the reference type $REF(ST)$.
 - i) <subtable clause> shall not be specified.
 - ii) <table scope> shall not be specified.
 - iii) If SYSTEM GENERATED is specified, then RST shall have a system-defined representation.
 - iv) If USER GENERATED is specified, then RST shall have a user-defined representation.
 - v) If DERIVED is specified, then RST shall have a derived representation.
 - vi) If RST has a derived representation, then let m be the number of attributes included in the list of attributes of the derived representation of RST and let A_i , $1 \text{ (one)} \leq i \leq m$, be those attributes.
 - 1) TD shall contain a <table constraint definition> that specifies a <unique constraint definition> UCD whose <unique column list> contains the attribute names of A_1, A_2, \dots, A_m in that order.

- 2) If *UCD* does not specify PRIMARY KEY, then for every attribute A_i , $1 \text{ (one)} \leq i \leq m$, *TD* shall contain a <column options> CO_i with a <column name> that is equivalent to the <attribute name> of A_i and with a <column constraint definition> that specifies NOT NULL.

- vii) Let CD_0 be the <column definition>:

CN_0 RST SCOPE(*TN*) UNIQUE NOT NULL

where CN_0 denotes the <self-referencing column name> simply contained in *SRCS*.

- d) If <subtable clause> is specified, then:

- i) The <table name> contained in the <subtable clause> identifies the *direct supertable* of *T*, which shall be a base table. *T* is called a *direct subtable* of the direct supertable of *T*.
- ii) *ST* shall be a direct subtype of the structured type of the direct supertable of *T*.
- iii) The SQL-schema identified by the explicit or implicit <schema name> of the <table name> of *T* shall include the descriptor of the direct supertable of *T*.
- iv) The subtable family of *T* shall not include a member, other than *T* itself, whose associated structured type is *ST*.
- v) *TD* shall not contain a <table constraint definition> that specifies PRIMARY KEY.
- vi) Let the term *inherited column* of *T* refer to a column of *T* that corresponds to an inherited attribute of *ST*. For every such inherited attribute *IA*, there is a column *CA* of the direct supertable of *T* such that the <column name> of *CA* is equivalent to the <attribute name> of *IA*. *CA* is called the *direct supercolumn* of *IA* in the direct supertable of *T*.
- vii) Let CD_0 be the <column definition>:

CN_0 RST SCOPE(*TN*) UNIQUE NOT NULL

where CN_0 denotes the <self-referencing column name> simply contained in *SRCS*.

- e) 13 Let the term *originally-defined column* of *T* refer to a column of *T* that corresponds to an originally-defined attribute of *ST*.
- f) Let *n* be the number of attributes of *ST*. Let AD_i , $1 \text{ (one)} \leq i \leq n$, be the attribute descriptors included in the data type descriptor of *ST* and let CD_i be the <column definition> CN_i DT_i DC_i , where:
- i) CN_i is the attribute name included in AD_i .
 - ii) DT_i is some <data type> that, under the General Rules of [Subclause 6.1](#), “<data type>”, would result in the creation of the data type descriptor included in AD_i .
 - iii) If AD_i describes an inherited attribute *IA*, then

Case:

 - 1) If the column descriptor of the direct supercolumn of *IA* includes a default value, then DC_i is some <default clause> whose <default option> denotes this default value.
 - 2) Otherwise, DC_i is the zero-length string.

- iv) If AD_i describes an originally-defined attribute OA , then

Case:

- 1) If AD_i includes a default value, then DC_i is some <default clause> whose <default option> denotes this default value.
- 2) Otherwise, DC_i is the zero-length string.

- g) If <typed table element list> $TTEL$ is specified and <column options> CO is specified, then:

- i) The <column name> CN simply contained in CO shall be equivalent to the <column name> CN_j specified in some <column definition> CD_j and shall refer to an originally-defined column of T .
- ii) CN shall not be equivalent to the <column name> simply contained in any other <column options> contained in $TTEL$.
- iii) 09 A <column option list> shall immediately contain either a <scope clause> or a <default clause>, or at least one <column constraint definition>.
- iv) If CO specifies a <scope clause> SC , then DT_j shall be a <reference type> RT . If RT contains a <scope clause>, then that <scope clause> is replaced by SC ; otherwise, RT is replaced by $RT\ SC$.

NOTE 276 — Changes to the scope of a column of a typed table do not affect the scope defined for the underlying attribute. Such an attribute scope serves as a kind of default for the column's scope, at the time the typed table is defined, and is not restored if a column's scope is dropped.

- v) If CO specifies a <default clause> DC , then DC_j is replaced by DC in CD_j .
 - vi) 09 If CO specifies a non-empty list $CCDL$ of <column constraint definition>s, then CD_j is replaced by $CD_j\ CCDL$.
 - vii) CO is deleted from $TTEL$.
- h) T is a referenceable table.
- i) If $TTEL$ is empty, then let TEL be a <table element list> of the form

$$CD_0, \dots, CD_n$$

Otherwise, then let TEL be a <table element list> of the form

$$CD_0, \dots, CD_n\ TTEL$$

- 9) If ON COMMIT is specified, then TEMPORARY shall be specified.
- 10) If TEMPORARY is specified and ON COMMIT is not specified, then ON COMMIT DELETE ROWS is implicit.
- 11) Every referenceable table referenced by a <scope clause> contained in a <column definition> or <column options> contained in TD shall be

Case:

- a) If TD specifies no <table scope>, then a persistent base table.

- b) If *TD* specifies GLOBAL TEMPORARY, then a global temporary table.
 - c) If *TD* specifies LOCAL TEMPORARY, then a created local temporary table.
- 12) At most one <table element> shall be a <column definition> that contains an <identity column specification>.
- 13) The <table name> is a range variable whose scope is the <table definition>, excluding the <as subquery clause>.

Access Rules

- 1) If a <table definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.
- 2) If a <like clause> is contained in a <table definition>, then the applicable privileges for *A* shall include SELECT privilege on the table identified in the <like clause>.
- 3) *A* shall have in its applicable privileges the UNDER privilege on the <supertable name> specified in <subtable clause>.
- 4) If “OF <path-resolved user-defined type name>” is specified, then the applicable privileges for *A* shall include USAGE on *ST*.

General Rules

- 1) A <table definition> defines either a persistent base table, a global temporary table or a created local temporary table. If GLOBAL is specified, then a global temporary table is defined. If LOCAL is specified, then a created local temporary table is defined. Otherwise, a persistent base table is defined.
- 2) 09 The degree of *T* is initially set to 0 (zero); the General Rules of [Subclause 11.4, “<column definition>”](#), specify the degree of *T* during the definition of the columns of *T*.
- 3) If <path-resolved user-defined type name> is specified, then:
 - a) Let *R* be the structured type identified by the <user-defined type name> simply contained in <path-resolved user-defined type name>.
 - b) *R* is the structured type associated with *T*.
- 4) A table descriptor *TDS* is created that describes *T*. *TDS* includes:
 - a) The table name *TN*.
 - b) The column descriptors of every column of *T*, according to the Syntax Rules and General Rules of [Subclause 11.4, “<column definition>”](#), applied to the <column definition>s contained in *TEL*, in the order in which they were specified.
 - c) If <typed table clause> is specified, then:
 - i) An indication that the table is a referenceable table.
 - ii) An indication that the column at ordinal position 1 (one) is the self-referencing column of *T*. The column descriptor included in *TDS* that describes that column is marked as identifying a self-referencing column.

- iii) If *RST* has a system-defined representation, then an indication that the self-referencing column is a system-generated self-referencing column.
 - iv) If *RST* has a derived representation, then an indication that the self-referencing column is a derived self-referencing column.
 - v) If *RST* has a user-defined representation, then an indication that the self-referencing column is a user-generated self-referencing column.
- d) The table constraint descriptors specified by each <table constraint definition> contained in *TEL*.
- e) If a <path-resolved user-defined type name> is specified, then the user-defined type name of *R*.
- f) If <subtable clause> is specified, then the table name of the direct supertable of *T* contained in the <subtable clause>.
- g) A non-empty set of functional dependencies, according to the rules given in Subclause 4.18, “Functional dependencies”.
- h) A non-empty set of candidate keys.
- i) A preferred candidate key, which may or may not be additionally designated the primary key, according to the Rules in Subclause 4.18, “Functional dependencies”.
- j) An indication of whether the table is a persistent base table, a global temporary table, a created local temporary table, or a declared local temporary table.
- k) If TEMPORARY is specified, then
- Case:
- i) If ON COMMIT PRESERVE ROWS is specified, then the table descriptor includes an indication that ON COMMIT PRESERVE ROWS is specified.
 - ii) Otherwise, the table descriptor includes an indication that ON COMMIT DELETE ROWS is specified or implied.
- l) Case:
- i) If <typed table clause> is not specified, then an indication that *T* is insertable-into and updatable.
 - ii) Otherwise,
- Case:
- 1) If the data type descriptor of *R* indicates that *R* is instantiable, then an indication that *T* is insertable-into and updatable.
 - 2) Otherwise, an indication that *T* is not insertable-into and not updatable.
- 5) In the descriptor of each direct supertable of *T*, *TN* is added to the end of the list of direct subtables.
- 6) If <subtable clause> is specified, then a set of privilege descriptors is created that defines the privileges SELECT, UPDATE, and REFERENCES for every inherited column of this table to the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of the <table name> of the direct supertable from which that column was inherited. These privileges are grantable. The grantor for each of these privilege descriptors is set to the special grantor value “_SYSTEM”.

11.3 <table definition>

- 7) A set of privilege descriptors is created that define the privileges INSERT, SELECT, UPDATE, DELETE, TRIGGER, and REFERENCES on this table and SELECT, INSERT, UPDATE, and REFERENCES for every <column definition> in the table definition. If OF <path-resolved user-defined type name> is specified, then a table/method privilege descriptor is created on this table for every method of the structured type identified by the <path-resolved user-defined type name> and the table SELECT privilege has the WITH HIERARCHY OPTION. These privileges are grantable.

The grantor for each of these privilege descriptors is set to the special grantor value “_SYSTEM”. The grantee is <authorization identifier> A.

- 8) If <subtable clause> is specified, then let *ST* be the set of supertables of *T*. Let *PDS* be the set of privilege descriptors that defined SELECT WITH HIERARCHY OPTION privilege on a table in *ST*. For every privilege descriptor in *PDS*, with grantee *G*, grantor *A*,

Case:

- a) If the privilege is grantable, then let *WGO* be “WITH GRANT OPTION”.
- b) Otherwise, let *WGO* be a zero-length string.

The following <grant statement> is effectively executed without further Access Rule checking:

```
GRANT SELECT ON T TO G WGO FROM A
```

- 9) The row type *RT* of the table *T* defined by the <table definition> is the set of pairs (<field name>, <data type>) where <field name> is the name of a column *C* of *T* and <data type> is the declared type of *C*. This set of pairs contains one pair for each column of *T*, in the order of their ordinal position in *T*.
- 10) If <as subquery clause> is specified and WITH DATA is specified, then let *QE* be the <query expression> simply contained in *ASQ*. The following <insert statement> is effectively executed without further Access Rule checking:

```
INSERT INTO TN QE
```

Conformance Rules

- 1) Without Feature T171, “LIKE clause in table definition”, conforming SQL language shall not contain a <like clause>.
- 2) Without Feature F531, “Temporary tables”, conforming SQL language shall not contain a <table scope> and shall not reference any global or local temporary table.
- 3) Without Feature S051, “Create table of type”, conforming SQL language shall not contain “OF <path-resolved user-defined type name>”.
- 4) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <column option list> that contains a <scope clause>.
- 5) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain <reference generation> that does not contain SYSTEM GENERATED.
- 6) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <subtable clause>.
- 7) Without Feature T172, “AS subquery clause in table definition”, conforming SQL language shall not contain an <as subquery clause>.

- 8) Without Feature T173, “Extended LIKE clause in table definition”, a <like clause> shall not contain <like options>.

11.4 <column definition>

*This Subclause is modified by Subclause 9.4, “<column definition>”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 12.1, “<column definition>”, in ISO/IEC 9075-14.*

Function

Define a column of a base table.

Format

```
<column definition> ::=
    <column name> [ <data type or domain name> ]
        [ <default clause> | <identity column specification> | <generation clause> ]
        [ <column constraint definition>... ]
        [ <collate clause> ]

<data type or domain name> ::=
    <data type>
    | <domain name>

<column constraint definition> ::=
    [ <constraint name definition> ] <column constraint> [ <constraint characteristics> ]

<column constraint> ::=
    NOT NULL
    | <unique specification>
    | <references specification>
    | <check constraint definition>

<identity column specification> ::=
    GENERATED { ALWAYS | BY DEFAULT } AS IDENTITY
        [ <left paren> <common sequence generator options> <right paren> ]

<generation clause> ::=
    <generation rule> AS <generation expression>

<generation rule> ::=
    GENERATED ALWAYS

14 <generation expression> ::=
    <left paren> <value expression> <right paren>
```

Syntax Rules

- 1) 04 The <column definition> shall not contain a <host parameter specification>, an <SQL parameter reference>, a <dynamic parameter specification>, or an <embedded variable specification>.
- 2) Case:
 - a) If the <column definition> is contained in a <table definition>, then let *T* be the table defined by that <table definition>.

- b) If the <column definition> is contained in a <temporary table declaration>, then let *T* be the table declared by that <temporary table declaration>.
- c) If the <column definition> is contained in an <alter table statement>, then let *T* be the table identified in the containing <alter table statement>.

The <column name> in the <column definition> shall not be equivalent to the <column name> of any other column of *T*.

- 3) Let *A* be the <authorization identifier> that owns *T*.
- 4) Let *C* be the <column name> of the <column definition>.
- 5) <data type or domain name> shall unambiguously reference either a <data type> or a <domain name>.
- 6) If <domain name> is specified, then let *D* be the domain identified by the <domain name>.
- 7) If <generation clause> *GC* is specified, then:
 - a) Let *GE* be the <generation expression> contained in *GC*.
 - b) *C* is a generated column.
 - c) Every <column reference> contained in *GE* shall reference a base column of *T*.
 - d) *GE* shall be deterministic.
 - e) *GE* shall not contain a <routine invocation> whose subject routine possibly reads SQL-data.
 - f) *GE* shall not contain a <query expression>.
- 8) If <generation clause> is omitted, then either <data type> or <domain name> shall be specified.
- 9) Case:
 - a) If <column definition> immediately contains <domain name>, then it shall not also immediately contain <collate clause>.
 - b) Otherwise, <collate clause> shall not be both specified in <data type> and immediately contained in <column definition>. If <collate clause> is immediately contained in <column definition>, then it is equivalent to specifying an equivalent <collate clause> in <data type>.
- 10) The declared type of the column is
Case:
 - a) If <data type> is specified, then that data type. If <generation clause> is also specified, then the declared type of <generation expression> shall be assignable to the declared type of the column.
 - b) If <domain name> is specified, then the declared type of *D*. If <generation clause> is also specified, then the declared type of <generation expression> shall be assignable to the declared type of the column.
 - c) If <generation clause> is specified, then the declared type of *GE*.
- 11) If a <data type> is specified, then:
 - a) Let *DT* be the <data type>.
 - b) If *DT* specifies CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT and does not specify a <character set specification>, then the <character set specification> specified or

implicit in the <schema character set specification> of the <schema definition> that created the schema identified by the <schema name> immediately contained in the <table name> of the containing <table definition> or <alter table statement> is implicit.

12) If <identity column specification> *ICS* is specified, then:

- a) Case:
 - i) If the declared type of the column being defined is a distinct type *DIST*, then the source type of *DIST* shall be exact numeric with scale 0 (zero). Let *ICT* be the source type of *DIST*.
 - ii) Otherwise, the declared type of the column being defined shall be exact numeric with scale 0 (zero). Let *ICT* be the declared type of the column being defined.
- b) Let *SGO* be the <common sequence generator options>.
- c) The Syntax Rules of Subclause 9.24, “Creation of a sequence generator”, are applied with *SGO* as *OPTIONS* and *ICT* as *DATA TYPE*.
- d) The <column constraint definition> NOT NULL NOT DEFERRABLE is implicit.

13) If a <column constraint definition> is specified, then let *CND* be the <constraint name definition> if one is specified and let *CND* be a zero-length string otherwise; let *CA* be the <constraint characteristics> if specified and let *CA* be a zero-length string otherwise. The <column constraint definition> is equivalent to a <table constraint definition> as follows.

Case:

- a) If a <column constraint definition> is specified that contains the <column constraint> NOT NULL, then it is equivalent to the following <table constraint definition>:

CND CHECK (*C* IS NOT NULL) *CA*

- b) If a <column constraint definition> is specified that contains a <unique specification> *US*, then it is equivalent to the following <table constraint definition>:

CND *US* (*C*) *CA*

NOTE 277 — The <unique specification> is defined in Subclause 11.7, “<unique constraint definition>”.

- c) If a <column constraint definition> is specified that contains a <references specification> *RS*, then it is equivalent to the following <table constraint definition>:

CND FOREIGN KEY (*C*) *RS* *CA*

NOTE 278 — The <references specification> is defined in Subclause 11.8, “<referential constraint definition>”.

- d) If a <column constraint definition> is specified that contains a <check constraint definition> *CCD*, then it is equivalent to the following <table constraint definition>:

CND *CCD* *CA*

Each column reference directly contained in the <search condition> shall reference column *C*.

14) The schema identified by the explicit or implicit qualifier of the <domain name> shall include the descriptor of *D*.

Access Rules

- 1) If <domain name> is specified, then the applicable privileges for *A* shall include USAGE on *D*.

General Rules

- 1) A <column definition> defines a column in a table.
- 2) If the <column definition> specifies <data type>, then a data type descriptor is created that describes the declared type of the column being defined.
- 3) The degree of the table *T* being defined in the containing <table definition> or <temporary table declaration>, or being altered by the containing <alter table statement> is increased by 1 (one).
- 4) A column descriptor is created that describes the column being defined. The column descriptor includes:
 - a) *C*, the name of the column.
 - b) Case:
 - i) If the <column definition> specifies a <data type> or a <generation clause>, then the data type descriptor of the declared type of the column.
 - ii) Otherwise, the <domain name> that identifies the domain of the column.
 - c) The ordinal position of the column, which is equal to the degree of *T*.
 - d) The nullability characteristic of the column, determined according to the rules in [Subclause 4.13, “Columns, fields, and attributes”](#).

NOTE 279 — Both <column constraint definition>s and <table constraint definition>s shall be analyzed to determine the nullability characteristics of all columns.

- e) If <default clause> is specified, then the <default option>.
- f) If <identity column specification> is specified, then:
 - i) An indication that the column is an identity column.
 - ii) If ALWAYS is specified, then an indication that values are always generated.
 - iii) If BY DEFAULT is specified, then an indication that values are generated by default.
 - iv) The descriptor of the sequence generator *SG* resulting from application of the General Rules of [Subclause 9.24, “Creation of a sequence generator”](#), with *SGO* as *OPTIONS* and *ICT* as *DATA TYPE*.
 - v) The next available value of *SG* as the start value.
- g) Case:
 - i) If the <column definition> specifies a <generation clause>, then an indication that the column is “ALWAYS” generated and *GE*.
 - ii) Otherwise, an indication that the column is “NEVER” generated.
- h) An indication that the column is updatable.

- 5) If <domain name> is specified, then, for every domain constraint descriptor *DCD* included in the domain descriptor of *D*, let *DSC* be the template <search condition> included in *DCD*. Let *CSC* be a copy of *DSC* in which every instance of the <general value specification> *VALUE* is replaced by *C*. A domain constraint usage descriptor is created and added to the set of domain constraint usage descriptors included in *DCD*. The domain constraint usage descriptor created includes:
- a) The name of the applicable column.
 - b) The constraint name included in *DCD*, as the name of the applicable domain constraint.
 - c) The indication included in *DCD* as to whether the constraint is deferrable.
 - d) The initial constraint mode included in *DCD*.
 - e) The applicable <search condition>

```
( SELECT EVERY ( CSC )  
  FROM T )
```

NOTE 280 — This is a <scalar subquery> of declared type BOOLEAN.

Conformance Rules

- 1) Without Feature F692, “Extended collation support”, conforming SQL language shall not contain a <column definition> that immediately contains a <collate clause>.
- 2) Without Feature T174, “Identity columns”, conforming SQL language shall not contain an <identity column specification>.
- 3) Without Feature T175, “Generated columns”, conforming SQL language shall not contain a <generation clause>.

11.5 <default clause>

This Subclause is modified by Subclause 9.5, “<default clause>”, in ISO/IEC 9075-4.

Function

Specify the default for a column, domain, or attribute.

Format

```
<default clause> ::=
    DEFAULT <default option>

<default option> ::=
    <literal>
  | <datetime value function>
  | USER
  | CURRENT_USER
  | CURRENT_ROLE
  | SESSION_USER
  | SYSTEM_USER
  | CURRENT_CATALOG
  | CURRENT_SCHEMA
  | CURRENT_PATH
  | <implicitly typed value specification>
```

Syntax Rules

- 1) 04 The subject data type of a <default clause> is the data type specified in the descriptor identified by the containing <column definition>, <domain definition>, <attribute definition>, <alter column definition>, or <alter domain statement>.
- 2) If USER is specified, then CURRENT_USER is implicit.
- 3) Case:
 - a) If the subject data type of the <default clause> is a user-defined type, a reference type, or a row type, then <default option> shall specify <null specification>.
 - b) If the subject data type of the <default clause> is a collection type, then <default option> shall specify <implicitly typed value specification>. If the <default option> specifies an <empty specification> that specifies ARRAY, then the subject data type shall be an array type. If the <default option> specifies an <empty specification> that specifies MULTISSET, then the subject data type shall be a multiset type.
- 4) Case:
 - a) If a <literal> is specified, then
Case:
 - i) If the subject data type is character string, then the <literal> shall be a <character string literal>. If the length of the subject data type is fixed, then the length in characters of the <character string literal> shall not be greater than the length of the subject data type. If the length of the

subject data type is variable, then the length in characters of the <character string literal> shall not be greater than the maximum length of the subject data type. The <literal> shall have the same character repertoire as the subject data type.

- ii) If the subject data type is binary string, then the <literal> shall be a <binary string literal>. If the length of the subject data type is fixed, then the length in octets of the <binary string literal> shall not be greater than the length of the subject data type. If the length of the subject data type is variable, then the length in octets of the <binary string literal> shall not be greater than the maximum length of the subject data type.
 - iii) If the subject data type is exact numeric, then the <literal> shall be a <signed numeric literal> that simply contains an <exact numeric literal>. There shall be a representation of the value of the <literal> in the subject data type that does not lose any significant digits.
 - iv) If the subject data type is approximate numeric, then the <literal> shall be a <signed numeric literal>.
 - v) If the subject data type is datetime, then the <literal> shall be a <datetime literal> with the same primary datetime fields and the same time zone datetime fields as the subject data type. If SECOND is one of these fields, then the fractional seconds precision of the <datetime literal> shall be less than or equal to the fractional seconds precision of the subject data type.
 - vi) If the subject data type is interval, then the <literal> shall be an <interval literal> and shall contain the same <interval qualifier> as the subject data type.
 - vii) If the subject data type is boolean, then the <literal> shall be a <boolean literal>.
- b) If CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, or CURRENT_SCHEMA is specified, then the subject data type shall be character string with character set SQL_IDENTIFIER. If the length of the subject data type is fixed, then its length shall not be less than 128 characters. If the length of the subject data type is variable, then its maximum length shall not be less than 128 characters.
 - c) If CURRENT_PATH is specified, then the subject data type shall be character string with character set SQL_IDENTIFIER. If the length of the subject data type is fixed, then its length shall not be less than 1031 characters. If the length of the subject data type is variable, then its maximum length shall not be less than 1031 characters.
 - d) If <datetime value function> is specified, then the subject data type shall be datetime with the same declared datetime data type of the <datetime value function>.
 - e) If <empty specification> is specified, then the subject data type shall be a collection type.

Access Rules

None.

General Rules

- 1) The default value inserted in the column descriptor, if the <default clause> is to apply to a column, or in the domain descriptor, if the <default clause> is to apply to a domain, or in the attribute descriptor, if the <default clause> is to apply to an attribute, is the <default option>.

2) The value specified by a <default option> is

Case:

a) If the <default option> contains a <literal>, then

Case:

- i) If the subject data type is numeric, then the numeric value of the <literal>.
 - ii) If the subject data type is character string with variable length, then the value of the <literal>.
 - iii) If the subject data type is character string with fixed length, then the value of the <literal>, extended as necessary on the right with <space>s to the length in characters of the subject data type.
 - iv) If the subject data type is fixed-length binary string, then the value of the <literal>, extended as necessary to the right with X'00's to the length in octets of the subject data type.
 - v) If the subject data type is variable-length binary string, then the value of the <literal>.
 - vi) If the subject data type is datetime or interval, then the value of the <literal>.
 - vii) If the subject data type is boolean, then the value of the <literal>.
- b) If the <default option> specifies CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, CURRENT_SCHEMA, or CURRENT_PATH, then

Case:

- i) If the subject data type is character string with variable length, then the value obtained by an evaluation of CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, CURRENT_SCHEMA, or CURRENT_PATH at the time that the default value is required.
 - ii) If the subject data type is character string with fixed length, then the value obtained by an evaluation of CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, CURRENT_SCHEMA, or CURRENT_PATH at the time that the default value is required, extended as necessary on the right with <space>s to the length in characters of the subject data type.
- c) If the <default option> contains a <datetime value function>, then the value of an evaluation of the <datetime value function> at the time that the default value is required.
- d) If the <default option> specifies <empty specification>, then an empty collection.

3) When a site *S* is set to its default value,

Case:

- a) If the descriptor of *S* indicates that it represents a column of which some underlying column is an identity column or a generated column, then *S* is marked as *unassigned*.

NOTE 281 — The notion of a site being unassigned is only for definitional purposes in this International Standard. It is not a state that can persist so as to be visible in SQL-data. The treatment of unassigned sites is given in Subclause 15.10, “Effect of inserting tables into base tables”, and Subclause 15.13, “Effect of replacing rows in base tables”.

- b) If the data descriptor for the site includes a <default option>, then *S* is set to the value specified by that <default option>.

11.5 <default clause>

- c) If the data descriptor for the site includes a <domain name> that identifies a domain descriptor that includes a <default option>, then *S* is set to the value specified by that <default option>.
- d) If the default value is for a column *C* of a candidate row for insertion into or update of a derived table *DT* and *C* has a single counterpart column *CC* in a leaf generally underlying table of *DT*, then *S* is set to the default value of *CC*, which is obtained by applying the General Rules of this Subclause.
- e) Otherwise, *S* is set to the null value.

NOTE 282 — If <default option> specifies CURRENT_USER, CURRENT_ROLE, SESSION_USER, SYSTEM_USER, CURRENT_CATALOG, CURRENT_SCHEMA, or CURRENT_PATH, then the “value in the column descriptor” will effectively be the text of the <default option>, whose evaluation occurs at the time that the default value is required.

- 4) If the <default clause> is contained in an <SQL schema statement> and character representation of the <default option> cannot be represented in the Information Schema without truncation, then a completion condition is raised: *warning — default value too long for information schema*.

NOTE 283 — The Information Schema is defined in [ISO9075-11].

Conformance Rules

- 1) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <default option> that contains CURRENT_PATH.
- 2) Without Feature F321, “User authorization”, conforming SQL language shall not contain a <default option> that contains CURRENT_USER, SESSION_USER, or SYSTEM_USER.

NOTE 284 — Although CURRENT_USER and USER are semantically the same, without Feature F321, “User authorization”, CURRENT_USER shall be specified as USER.

- 3) Without Feature T332, “Extended roles”, conforming SQL language shall not contain a <default option> that contains CURRENT_ROLE.
- 4) Without Feature F762, “CURRENT_CATALOG”, conforming SQL language shall not contain a <default option> that contains CURRENT_CATALOG.
- 5) Without Feature F763, “CURRENT_SCHEMA”, conforming SQL language shall not contain a <default option> that contains CURRENT_SCHEMA.

11.6 <table constraint definition>

Function

Specify an integrity constraint.

Format

```
<table constraint definition> ::=  
  [ <constraint name definition> ] <table constraint>  
  [ <constraint characteristics> ]  
  
<table constraint> ::=  
  <unique constraint definition>  
  | <referential constraint definition>  
  | <check constraint definition>
```

Syntax Rules

- 1) If <constraint characteristics> is not specified, then INITIALLY IMMEDIATE NOT DEFERRABLE is implicit.
- 2) If <constraint name definition> is specified and its <constraint name> contains a <schema name>, then that <schema name> shall be equivalent to the explicit or implicit <schema name> of the <table name> of the table identified by the containing <table definition> or <alter table statement>.
- 3) If <constraint name definition> is not specified, then a <constraint name definition> that contains an implementation-dependent <constraint name> is implicit. The assigned <constraint name> shall obey the Syntax Rules of an explicit <constraint name>.
- 4) Let *S* be the schema identified by the explicit or implicit <schema name> of the <constraint name>. *S* shall not include a constraint descriptor whose constraint name is <constraint name>.

Access Rules

None.

General Rules

- 1) A <table constraint definition> defines a table constraint.
- 2) A table constraint descriptor is created that describes the table constraint being defined. The table constraint descriptor includes:
 - a) The <constraint name> contained in the explicit or implicit <constraint name definition>.
 - b) An indication of whether the constraint is deferrable or not deferrable.
 - c) An indication of whether the initial constraint mode of the constraint is deferred or immediate.

11.6 <table constraint definition>

- d) The applicable <search condition> as specified in the General Rules applicable to the particular kind of <table constraint> contained in the <table constraint definition>, and any additional contents specified by those General Rules.
- 3) If the <table constraint> is a <check constraint definition>, then let *SC* be the <search condition> immediately contained in the <check constraint definition> and let *T* be the table name included in the corresponding table constraint descriptor; the applicable <search condition> included in the table constraint descriptor is:

```
( SELECT EVERY ( SC )  
  FROM T )
```

NOTE 285 — This is a <scalar subquery> of declared type BOOLEAN.

Conformance Rules

None.

11.7 <unique constraint definition>

This Subclause is modified by *Subclause 11.4*, “<unique constraint definition>”, in ISO/IEC 9075-9.

Function

Specify a uniqueness constraint for a table.

Format

```
<unique constraint definition> ::=  
    <unique specification> <left paren> <unique column list> <right paren>  
    | UNIQUE ( VALUE )  
  
<unique specification> ::=  
    UNIQUE  
    | PRIMARY KEY  
  
<unique column list> ::=  
    <column name list>
```

Syntax Rules

- 1) 09 Each column identified by a <column name> in the <unique column list> is an operand of a grouping operation. The Syntax Rules of *Subclause 9.10*, “Grouping operations”, apply.
- 2) Let *T* be the table identified by the containing <table definition> or <alter table statement>. Let *TN* be the <table name> of *T*.
- 3) If <unique column list> *UCL* is specified, then
 - a) Each <column name> in the <unique column list> shall identify a column of *T*, and the same column shall not be identified more than once.
 - b) The set of columns in the <unique column list> shall be distinct from the unique columns of any other unique constraint descriptor that is included in the base table descriptor of *T*.
 - c) Case:
 - i) If the <unique specification> specifies PRIMARY KEY, then let *SC* be the <search condition>:

```
UNIQUE ( SELECT UCL  
          FROM TN )  
AND  
( SELECT EVERY ( UCL ) IS NOT NULL )  
  FROM TN )
```

NOTE 286 — The second operand of AND in this expression is a <scalar subquery> of declared type BOOLEAN.
 - ii) Otherwise, let *SC* be the <search condition>:

```
UNIQUE ( SELECT UCL FROM TN )
```
- 4) If UNIQUE (VALUE) is specified, then let *SC* be the <search condition>:

11.7 <unique constraint definition>

UNIQUE (SELECT TN.* FROM TN)

- 5) If the <unique specification> specifies PRIMARY KEY, then for each <column name> in the explicit or implicit <unique column list> for which NOT NULL is not specified, NOT NULL is implicit in the <column definition>.
- 6) A <table definition> shall specify at most one implicit or explicit <unique constraint definition> that specifies PRIMARY KEY.
- 7) If a <unique constraint definition> that specifies PRIMARY KEY is contained in an <add table constraint definition>, then the table identified by the <table name> immediately contained in the containing <alter table statement> shall not have a unique constraint that was defined by a <unique constraint definition> that specified PRIMARY KEY.

Access Rules

None.

General Rules

- 1) A <unique constraint definition> defines a unique constraint.
- 2) The applicable <search condition> included in the table constraint descriptor created as a result of execution of the containing <table constraint definition> is SC. This descriptor additionally includes:
 - a) The <unique specification>, indicating whether the constraint is defined with PRIMARY KEY or UNIQUE.
 - b) The names of the unique columns specified in the <unique column list>.
- 3) The unique constraint is not satisfied if and only if

EXISTS (SELECT * FROM TN WHERE NOT (SC))

is True.

Conformance Rules

- 1) Without Feature S291, “Unique constraint on entire row”, conforming SQL language shall not contain UNIQUE(VALUE).
- 2) Without Feature T591, “UNIQUE constraints of possibly null columns”, in conforming SQL language, if UNIQUE is specified, then the <column definition> for each column whose <column name> is contained in the <unique column list> shall contain NOT NULL.

NOTE 287 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

11.8 <referential constraint definition>

Function

Specify a referential constraint.

Format

```
<referential constraint definition> ::=
    FOREIGN KEY <left paren> <referencing columns> <right paren>
        <references specification>

<references specification> ::=
    REFERENCES <referenced table and columns>
        [ MATCH <match type> ] [ <referential triggered action> ]

<match type> ::=
    FULL
    | PARTIAL
    | SIMPLE

<referencing columns> ::=
    <reference column list>

<referenced table and columns> ::=
    <table name> [ <left paren> <reference column list> <right paren> ]

<reference column list> ::=
    <column name list>

<referential triggered action> ::=
    <update rule> [ <delete rule> ]
    | <delete rule> [ <update rule> ]

<update rule> ::=
    ON UPDATE <referential action>

<delete rule> ::=
    ON DELETE <referential action>

<referential action> ::=
    CASCADE
    | SET NULL
    | SET DEFAULT
    | RESTRICT
    | NO ACTION
```

Syntax Rules

- 1) If <match type> is not specified, then SIMPLE is implicit.
- 2) Let *referencing table* be the table identified by the containing <table definition> or <alter table statement>. Let *referenced table* be the table identified by the <table name> in the <referenced table and columns>.

11.8 <referential constraint definition>

Let *referencing columns* be the column or columns identified by the <reference column list> in the <referencing columns> and let *referencing column* be one such column.

3) Case:

- a) If the <referenced table and columns> specifies a <reference column list>, then there shall be a one-to-one correspondence between the set of <column name>s contained in that <reference column list> and the set of <column name>s contained in the <unique column list> of a unique constraint of the referenced table such that corresponding <column name>s are equivalent. Let *referenced columns* be the column or columns identified by that <reference column list> and let *referenced column* be one such column. Each referenced column shall identify a column of the referenced table and the same column shall not be identified more than once.
 - b) If the <referenced table and columns> does not specify a <reference column list>, then the table descriptor of the referenced table shall include a unique constraint that specifies PRIMARY KEY. Let *referenced columns* be the column or columns identified by the unique columns in that unique constraint and let *referenced column* be one such column. The <referenced table and columns> shall be considered to implicitly specify a <reference column list> that is identical to that <unique column list>.
- 4) The table constraint descriptor describing the <unique constraint definition> whose <unique column list> identifies the referenced columns shall indicate that the unique constraint is not deferrable.
- 5) The referenced table shall be a base table.

Case:

- a) If the referencing table is a persistent base table, then the referenced table shall be a persistent base table.
 - b) If the referencing table is a global temporary table, then the referenced table shall be a global temporary table.
 - c) If the referencing table is a created local temporary table, then the referenced table shall be either a global temporary table or a created local temporary table.
 - d) If the referencing table is a declared local temporary table, then the referenced table shall be either a global temporary table, a created local temporary table or a declared local temporary table.
- 6) If the referenced table is a temporary table with ON COMMIT DELETE ROWS specified, then the referencing table shall specify ON COMMIT DELETE ROWS.
- 7) Each referencing column shall identify a column of the referencing table, and the same column shall not be identified more than once.
- 8) Each referencing column is an operand of a grouping operation. The Syntax Rules of [Subclause 9.10, "Grouping operations"](#), apply.
- 9) The <referencing columns> shall contain the same number of <column name>s as the <referenced table and columns>. The *i*-th column identified in the <referencing columns> corresponds to the *i*-th column identified in the <referenced table and columns>. The declared type of each referencing column shall be comparable to the declared type of the corresponding referenced column. There shall not be corresponding constituents of the declared type of a referencing column and the declared type of the corresponding referenced column such that one constituent is datetime with time zone and the other is datetime without time zone.

- 10) If a <referential constraint definition> does not specify any <update rule>, then an <update rule> with a <referential action> of NO ACTION is implicit.
- 11) If a <referential constraint definition> does not specify any <delete rule>, then a <delete rule> with a <referential action> of NO ACTION is implicit.
- 12) If any referencing column is a generated column, then:
 - a) <referential action> shall not specify SET NULL or SET DEFAULT.
 - b) <update rule> shall not specify ON UPDATE CASCADE.
- 13) Let T be the referenced table. The schema identified by the explicit or implicit qualifier of the <table name> shall include the descriptor of T .

Access Rules

- 1) The applicable privileges for the owner of T shall include REFERENCES for each referenced column.

General Rules

- 1) A <referential constraint definition> defines a referential constraint.
- 2) Let R_f be the referencing columns in the referencing table U and let R_t be the referenced columns in the referenced table T . Let MP be

Case:

- a) If SIMPLE is specified or implicit, then

R_f MATCH SIMPLE (SELECT R_t FROM T)

- b) If PARTIAL is specified, then

R_f MATCH PARTIAL (SELECT R_t FROM T)

- c) If FULL is specified, then

R_f MATCH FULL (SELECT R_t FROM T)

- 3) The applicable <search condition> included in the table constraint descriptor created as a result of executing the containing <table constraint definition> is:

(SELECT EVERY (MP)
FROM U)

NOTE 288 — This is a <scalar subquery> of declared type BOOLEAN.

This descriptor additionally includes:

- a) A list of the names of the referencing columns specified in the <referencing columns>, in an implementation-defined order.
- b) The name of the referenced table specified in the <referenced table and columns>.

11.8 <referential constraint definition>

- c) A list of the names of the referenced columns specified in the <referenced table and columns>, ordered such that each element is in the same position as that of its corresponding referencing column name in the list of names of referencing columns.
- d) The explicit or implicit <match type>.
- e) The <referential triggered action>, comprising the explicit or implicit <delete rule> and the explicit or implicit <update rule>.

Conformance Rules

- 1) Without Feature T191, “Referential action RESTRICT”, conforming SQL language shall not contain a <referential action> that contains RESTRICT.
- 2) Without Feature F741, “Referential MATCH types”, conforming SQL language shall not contain a <references specification> that contains MATCH.
- 3) Without Feature F191, “Referential delete actions”, conforming SQL language shall not contain a <delete rule>.
- 4) Without Feature F701, “Referential update actions”, conforming SQL language shall not contain an <update rule>.
- 5) Without Feature T201, “Comparable data types for referential constraints”, conforming SQL language shall not contain a <referencing columns> in which the data type of each referencing column is not the same as the data type of the corresponding referenced column.

NOTE 289 — The Conformance Rules of [Subclause 9.10](#), “[Grouping operations](#)”, also apply.

11.9 <check constraint definition>

*This Subclause is modified by Subclause 9.6, “<check constraint definition>”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 11.5, “<check constraint definition>”, in ISO/IEC 9075-9.
This Subclause is modified by Subclause 12.2, “<check constraint definition>”, in ISO/IEC 9075-14.*

Function

Specify a condition for the SQL-data.

Format

```
14 <check constraint definition> ::=  
CHECK <left paren> <search condition> <right paren>
```

Syntax Rules

- 1) 04 The <search condition> shall not contain a <host parameter specification>, an <SQL parameter reference>, a <dynamic parameter specification>, or an <embedded variable specification>.
- 2) The <search condition> shall not contain a <set function specification> that is not contained in a <query expression>.
- 3) If <check constraint definition> is contained in a <table definition> or <alter table statement>, then let *T* be the table identified by the containing <table definition> or <alter table statement>.

Case:

- a) If *T* is a persistent base table, or if the <check constraint definition> is contained in a <domain definition> or <alter domain statement>, then no <table reference> generally contained in the <search condition> shall reference a temporary table.
 - b) If *T* is a global temporary table, then no <table reference> generally contained in the <search condition> shall reference a table other than a global temporary table.
 - c) If *T* is a created local temporary table, then no <table reference> generally contained in the <search condition> shall reference a table other than either a global temporary table or a created local temporary table.
 - d) If *T* is a declared local temporary table, then no <table reference> generally contained in the <search condition> shall reference a persistent base table.
- 4) If the <check constraint definition> is contained in a <table definition> that defines a temporary table and specifies ON COMMIT PRESERVE ROWS or a <temporary table declaration> that specifies ON COMMIT PRESERVE ROWS, then the <search condition> shall not contain a reference to a temporary table defined by a <table definition> or a <temporary table declaration> that specifies ON COMMIT DELETE ROWS.
 - 5) The <search condition> shall simply contain a <boolean value expression> that is retrospectively deterministic.

NOTE 290 — “retrospectively deterministic” is defined in Subclause 6.34, “<boolean value expression>”.

11.9 <check constraint definition>

- 6) The <search condition> shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.
- 7) Let *A* be the <authorization identifier> that owns *T*.

Access Rules

None.

General Rules

- 1) A <check constraint definition> defines a check constraint.
- 2) If the character representation of the <search condition> cannot be represented in the Information Schema without truncation, then a completion condition is raised: *warning — search condition too long for information schema.*

NOTE 291 — The Information Schema is defined in [ISO9075-11].

Conformance Rules

- 1) Without Feature F671, “Subqueries in CHECK constraints”, conforming SQL language shall not contain a <search condition> contained in a <check constraint definition> that contains a <query expression>.
- 2) Without Feature F672, “Retrospective check constraints”, conforming SQL language shall not contain a <check constraint definition> that generally contains CURRENT_DATE, CURRENT_TIMESTAMP, or LOCALTIMESTAMP.

11.10 <alter table statement>

Function

Change the definition of a table.

Format

```
<alter table statement> ::=  
    ALTER TABLE <table name> <alter table action>  
  
<alter table action> ::=  
    <add column definition>  
    | <alter column definition>  
    | <drop column definition>  
    | <add table constraint definition>  
    | <drop table constraint definition>
```

Syntax Rules

- 1) Let T be the table identified by the <table name>.
- 2) The schema identified by the explicit or implicit schema name of the <table name> shall include the descriptor of T .
- 3) The scope of the <table name> is the entire <alter table statement>.
- 4) T shall be a base table.
- 5) T shall not be a declared local temporary table.

Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the <schema name> of the table identified by <table name>.

General Rules

- 1) The base table descriptor of T is modified as specified by <alter table action>.
- 2) If <add column definition> or <drop column definition> is specified, then the row type RT of T is the set of pairs (<field name>, <data type>) where <field name> is the name of a column C of T and <data type> is the declared type of C . This set of pairs contains one pair for each column of T in the order of their ordinal position in T .

Conformance Rules

None.

11.11 <add column definition>

Function

Add a column to a table.

Format

```
<add column definition> ::=
  ADD [ COLUMN ] <column definition>
```

Syntax Rules

- 1) Let *T* be the table identified by the <table name> immediately contained in the containing <alter table statement>.
- 2) *T* shall not be a referenceable table.
- 3) If <column definition> contains <identity column specification>, then the table descriptor of *T* shall not include a column descriptor of an identity column.

Access Rules

None.

General Rules

- 1) The column defined by the <column definition> is added to *T*.
- 2) Let *C* be the column added to *T*.

Case:

- a) If *C* is a generated column, then let *TN* be the <table name> immediately contained in the containing <alter table statement>, let *CN* be the <column name> immediately contained in <column definition>, and let *GE* be the generation expression included in the column descriptor of *C*. The following <update statement: searched> is executed without further Syntax Rule or Access Rule checking:

```
UPDATE TN SET CN = GE
```

- b) Otherwise, *C* is a base column.

Case:

- i) If *C* is an identity column, then for each row in *T* let *CS* be the site corresponding to *C* and let *NV* be the result of applying the General Rules of [Subclause 9.23](#), “[Generation of the next value of a sequence generator](#)”, with the sequence generator descriptor included in the column descriptor of *C* as *SEQUENCE*.

Case:

- 1) If the declared type of *C* is a distinct type *DIST*, then let *CNV* be *DIST(NV)*.
- 2) Otherwise, let *CNV* be *NV*.

The General Rules of Subclause 9.2, “Store assignment”, are applied with *CS* as *TARGET* and *CNV* as *VALUE*.

- ii) Otherwise, every value in *C* is the default value for *C*.

NOTE 292 — The default value of a column is defined in Subclause 11.5, “<default clause>”.

NOTE 293 — The addition of a column to a table has no effect on any existing <query expression> included in a view descriptor, <triggered action> included in a trigger descriptor, or <search condition> included in a constraint descriptor because any implicit column references in these descriptor elements are syntactically substituted by explicit column references under the Syntax Rules of Subclause 7.12, “<query specification>”. Furthermore, by implication (from the lack of any General Rules to the contrary), the meaning of a column reference is never retroactively changed by the addition of a column subsequent to the invocation of the <SQL schema statement> containing that column reference.

- 3) For every table privilege descriptor that specifies *T* and a privilege of SELECT, UPDATE, INSERT or REFERENCES, a new column privilege descriptor is created that specifies *T*, the same action, grantor, and grantee, and the same grantability, and specifies the <column name> of the <column definition>.
- 4) In all other respects, the specification of a <column definition> in an <alter table statement> has the same effect as specification of the <column definition> in the <table definition> for *T* would have had. In particular, the degree of *T* is increased by 1 (one) and the ordinal position of that column is equal to the new degree of *T* as specified in the General Rules of Subclause 11.4, “<column definition>”.

Conformance Rules

None.

11.12 <alter column definition>

Function

Change a column and its definition.

Format

```
<alter column definition> ::=
    ALTER [ COLUMN ] <column name> <alter column action>

<alter column action> ::=
    <set column default clause>
  | <drop column default clause>
  | <add column scope clause>
  | <drop column scope clause>
  | <alter column data type clause>
  | <alter identity column specification>
```

Syntax Rules

- 1) Let *T* be the table identified in the containing <alter table statement>.
- 2) Let *C* be the column identified by the <column name>.
- 3) *C* shall be a column of *T*.
- 4) If *C* is the self-referencing column of *T* or *C* is a generated column of *T*, then <alter column action> shall not contain <add column scope clause> or <drop column scope clause>.
- 5) If *C* is an identity column, then <alter column action> shall contain <alter identity column specification>.
- 6) If <alter identity column specification> is specified, then *C* shall be an identity column.

Access Rules

None.

General Rules

- 1) The column descriptor of *C* is modified as specified by <alter column action>.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <alter column definition>.

11.13 <set column default clause>

Function

Set the default clause for a column.

Format

```
<set column default clause> ::=  
    SET <default clause>
```

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) Let *C* be the column identified by the <column name> in the containing <alter column definition>.
- 2) The default value specified by the <default clause> is placed in the column descriptor of *C*.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <set column default clause>.

11.14 <drop column default clause>

Function

Drop the default clause from a column.

Format

```
<drop column default clause> ::=  
    DROP DEFAULT
```

Syntax Rules

- 1) Let *C* be the column identified by the <column name> in the containing <alter column definition>.
- 2) The descriptor of *C* shall include a default value.

Access Rules

None.

General Rules

- 1) The default value is removed from the column descriptor of *C*.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop column default clause>.

11.15 <add column scope clause>

Function

Add a non-empty scope for an existing column of data type REF in a base table.

Format

```
<add column scope clause> ::=  
  ADD <scope clause>
```

Syntax Rules

- 1) Let *C* be the column identified by the <column name> in the containing <alter column definition>. The declared type of *C* shall be some reference type. Let *RTD* be the reference type descriptor included in the descriptor of *C*.
- 2) Let *T* be the table identified by the <table name> in the containing <alter table statement>. If *T* is a referenceable table, then *C* shall be an originally-defined column of *T*.
- 3) *RTD* shall not include a scope.
- 4) Let *UDTN* be the name of the referenced type included in *RTD*.
- 5) The <table name> *STN* contained in the <scope clause> shall identify a referenceable table whose structured type is *UDTN*.

Access Rules

None.

General Rules

- 1) *STN* is included as the scope in the reference type descriptor included in the column descriptor of *C*.
- 2) For any proper subtable *PST* of *T*, let *PSC* be the column whose corresponding column in *T* is *C*. *STN* is included as the scope in the reference type descriptor included in the column descriptor of *PSC*.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <add column scope clause>.
- 2) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain an <add column scope clause>.

11.16 <drop column scope clause>

This Subclause is modified by Subclause 9.7, “<drop column scope clause>”, in ISO/IEC 9075-4.

Function

Drop the scope from an existing column of data type REF in a base table.

Format

```
<drop column scope clause> ::=  
  DROP SCOPE <drop behavior>
```

Syntax Rules

- 1) Let *C* be the column identified by the <column name> in the containing <alter column definition>. The declared type of *C* shall be some reference type whose reference type descriptor includes a scope.
- 2) Let *T* be the table identified by the <table name> in the containing <alter table statement>. If *T* is a referenceable table, then *C* shall be an originally-defined column of *T*.
- 3) Let *SC* be the set of columns consisting of *C* and, for every proper subtable of *T*, the column whose super-column is *C*.
- 4) An *impacted dereference operation* is a <dereference operation> whose <reference value expression> is a column reference that identifies a column in *SC*, a <method reference> whose <value expression primary> is a column reference that identifies a column in *SC*, or a <reference resolution> whose <reference value expression> is a column reference that identifies a column in *SC*.
- 5) If RESTRICT is specified, then no impacted dereference operation shall be contained in any of the following:
 - a) The SQL routine body of any routine descriptor.
 - b) The <query expression> of any view descriptor.
 - c) The <search condition> of any constraint descriptor.
 - d) 04 The triggered action of any trigger descriptor.

NOTE 294 — If CASCADE is specified, then such referencing objects will be dropped by the execution of the <SQL procedure statement>s specified in the General Rules of this Subclause.

Access Rules

None.

General Rules

- 1) 04 For every SQL-invoked routine *R* whose routine descriptor includes an SQL routine body that contains an impacted dereference operation, let *SN* be the <specific name> of *R*. The following <drop routine statement> is effectively executed for every *R* without further Access Rule checking:

DROP SPECIFIC ROUTINE *SN* CASCADE

- 2) For every view *V* whose view descriptor includes a <query expression> that contains an impacted dereference operation, let *VN* be the <table name> of *V*. The following <drop view statement> is effectively executed for every *V* without further Access Rule checking:

DROP VIEW *VN* CASCADE

- 3) For every assertion *A* whose assertion descriptor includes a <search condition> that contains an impacted dereference operation, let *AN* be the <constraint name> of *A*. The following <drop assertion statement> is effectively executed for every *A* without further Access Rule checking:

DROP ASSERTION *AN* CASCADE

- 4) 04 For every table check constraint *CC* whose table check constraint descriptor includes a <search condition> that contains an impacted dereference operation, let *CN* be the <constraint name> of *CC* and let *TN* be the <table name> of the table whose descriptor includes descriptor of *CC*. The following <alter table statement> is effectively executed for every *CC* without further Access Rule checking:

ALTER TABLE *TN* DROP CONSTRAINT *CN* CASCADE

- 5) The scope included in the reference type descriptor included in the column descriptor of every column in *SC* is made empty.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop column scope clause>.
- 2) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <drop column scope clause>.

11.17 <alter column data type clause>

Function

Change the declared type of a column.

Format

```
<alter column data type clause> ::=  
    SET DATA TYPE <data type>
```

Syntax Rules

- 1) Let *C* be the column identified by the <column name> in the containing <alter column definition>.
- 2) Let *T* be the table identified by the <table name> in the containing <alter table statement>.
- 3) *C* shall not be the self-referencing column of *T*.
- 4) *C* shall not be referenced in the <search condition> of any constraint descriptor other than a table constraint descriptor that contains references to no other column than *C* and that is included in the table descriptor of *T*.
- 5) *C* shall not be referenced in either an explicit trigger column list or a triggered action column set of any trigger descriptor.
- 6) *C* shall not be referenced in the SQL routine body of any routine descriptor.
- 7) *C* shall not be referenced in the <query expression> of any view descriptor.
- 8) *C* shall not be referenced in the generation expression of any column descriptor.
- 9) Let *D* be the data type specified by <data type>.
- 10) Let *DTC* be the declared type of *C*.
- 11) *D* shall be a predefined type.
- 12) *DTC* shall be a predefined type.
- 13) Case:
 - a) If *DTC* is a character string type, then:
 - i) *D* shall be a character string type.
 - ii) Let *LD* be the length or maximum length of *D*.
 - iii) Let *LDTC* be the length or maximum length of *DTC*.
 - iv) *LD* shall be greater than or equal to *LDTC*.
 - b) If *DTC* is a binary string type, then:
 - i) *D* shall be a binary string type.

- ii) Let LD be the length or maximum length of D .
- iii) Let $LDTC$ be the length or maximum length of DTC .
- iv) LD shall be greater than or equal to $LDTC$.
- c) If DTC is an exact numeric type, then:
 - i) D shall be an exact numeric type.
 - ii) Let PD and SD be the precision and scale, respectively, of D .
 - iii) Let $PDTC$ and $SDTC$ be the precision and scale, respectively, of DTC .
 - iv) PD shall be greater than or equal to $PDTC$.
 - v) SD shall be greater than or equal to $SDTC$.
 - vi) $(PD - SD)$ shall be greater than or equal to $(PDTC - SDTC)$.
- d) If DTC is an approximate numeric type, then:
 - i) D shall be an approximate numeric type.
 - ii) Let PD be the precision of D .
 - iii) Let $PDTC$ be the precision of DTC .
 - iv) PD shall be greater than or equal to $PDTC$.
- e) If DTC is DATE, then D shall be DATE.
- f) If DTC is a time type with time zone or a time type without time zone, then D shall be a time type with time zone or a time type without time zone, respectively.
 - i) Let PD be the value of the <time fractional seconds precision> of D .
 - ii) Let $PDTC$ be the value of the <time fractional seconds precision> of DTC .
 - iii) PD shall be greater than or equal to $PDTC$.
- g) If DTC is a timestamp type with time zone or a timestamp type without time zone, then D shall be a timestamp type with time zone or a timestamp type without time zone, respectively.
 - i) Let PD be the value of the <time fractional seconds precision> of D .
 - ii) Let $PDTC$ be the value of the <time fractional seconds precision> of DTC .
 - iii) PD shall be greater than or equal to $PDTC$.
- h) If DTC is a year-month interval type, then D shall be a year-month interval type.
- i) If DTC is a day-time interval type, then D shall be a day-time interval type.
- j) If DTC is BOOLEAN, then D shall be BOOLEAN.

Access Rules

None.

General Rules

- 1) The column descriptor of *C* is modified by replacing the existing data type descriptor with a data type descriptor describing *D*.

Conformance Rules

- 1) Without Feature F382, “Alter column data type”, conforming SQL language shall not contain an <alter column data type clause>.

11.18 <alter identity column specification>

Function

Change the options specified for an identity column.

Format

```
<alter identity column specification> ::=
  <alter identity column option>...

<alter identity column option> ::=
  <alter sequence generator restart option>
| SET <basic sequence generator option>
```

Syntax Rules

- 1) Let *SEQ* be the sequence generator descriptor included in the column descriptor identified by the <column name> in the containing <alter column definition>.
- 2) Let *OPT* be the character string formulated from <alter identity column specification> that conforms to the Format of <alter sequence generator options>.

NOTE 295 — *OPT* is formulated by removing all instances of the keyword SET from the string corresponding to from <alter identity column specification>.
- 3) The Syntax Rules of [Subclause 9.25, “Altering a sequence generator”](#), are applied with *OPT* as *OPTIONS* and *SEQ* as *SEQUENCE*.

Access Rules

None.

General Rules

- 1) The General Rules of [Subclause 9.25, “Altering a sequence generator”](#), are applied with *OPT* as *OPTIONS* and *SEQ* as *SEQUENCE*.

Conformance Rules

- 1) Without Feature T174, “Identity columns”, an <alter column definition> shall not contain an <alter identity column specification>.

11.19 <drop column definition>

*This Subclause is modified by Subclause 9.8, “<drop column definition>”, in ISO/IEC 9075-4.
This Subclause is modified by Subclause 11.6, “<drop column definition>”, in ISO/IEC 9075-9.*

Function

Destroy a column of a base table.

Format

```
<drop column definition> ::=
  DROP [ COLUMN ] <column name> <drop behavior>
```

Syntax Rules

- 1) Let *T* be the table identified by the <table name> in the containing <alter table statement> and let *TN* be the name of *T*.
- 2) Let *C* be the column identified by the <column name> *CN*.
- 3) *T* shall not be a referenceable table.
- 4) *C* shall be a column of *T* and *C* shall not be the only column of *T*.
- 5) 04 If RESTRICT is specified, then *C* shall not be referenced in any of the following:
 - a) The <query expression> of any view descriptor.
 - b) The <search condition> of any constraint descriptor other than a table constraint descriptor that contains references to no other column and that is included in the table descriptor of *T*.
 - c) The SQL routine body of any routine descriptor.
 - d) Either an explicit trigger column list or a triggered action column set of any trigger descriptor.
 - e) The generation expression of any column descriptor.

NOTE 296 — A <drop column definition> that does not specify CASCADE will fail if there are any references to that column resulting from the use of CORRESPONDING, NATURAL, SELECT * (except where contained in an <exists predicate>), a multi-column <unique constraint definition>, or REFERENCES without a <reference column list> in its <referenced table and columns>. Also note that a <drop column definition> whose <drop behavior> is RESTRICT that references a column that is referenced by a multi-column <unique constraint definition> will fail, while dropping a column that is referenced by a single-column <unique constraint definition> will succeed.

NOTE 297 — If CASCADE is specified, then any such dependent object will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.

Access Rules

None.

General Rules

- 1) Let *TR* be the trigger name of any trigger descriptor having an explicit trigger column list or a triggered action column set that contains *CN*. The following <drop trigger statement> is effectively executed without further Access Rule checking:

```
DROP TRIGGER TR
```

- 2) 09 Let *A* be the <authorization identifier> that owns *T*. The following <revoke statement> is effectively executed with a current authorization identifier of “_SYSTEM” and without further Access Rule checking:

```
REVOKE INSERT(CN) , UPDATE(CN) ,  
SELECT(CN) ,  
REFERENCES(CN) ON TABLE  
TN FROM A CASCADE
```

- 3) Let *GC* be any generated column of *T* in whose descriptor the generation expression contains a <column reference> that references *C*. The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE T DROP COLUMN GC CASCADE
```

- 4) If the column is not based on a domain, then its data type descriptor is destroyed; otherwise, all domain constraint usages that reference *C* are destroyed.
- 5) The data associated with *C* is destroyed.
- 6) The descriptor of *C* is removed from the descriptor of *T*.
- 7) The descriptor of *C* is destroyed.
- 8) The degree of *T* is reduced by 1 (one). The ordinal position of all columns having an ordinal position greater than the ordinal position of *C* is reduced by 1 (one).

Conformance Rules

- 1) Without Feature F033, “ALTER TABLE statement: DROP COLUMN clause”, conforming SQL language shall not contain a <drop column definition>.

11.20 <add table constraint definition>

Function

Add a constraint to a table.

Format

```
<add table constraint definition> ::=  
    ADD <table constraint definition>
```

Syntax Rules

- 1) Let T be the table identified by the <table name> contained in the containing <alter table statement>.
- 2) If PRIMARY KEY is specified, then T shall not have any proper supertable.

Access Rules

None.

General Rules

- 1) The table constraint descriptor for the <table constraint definition> is included in the table descriptor for T .
- 2) Let TC be the table constraint added to T . If TC causes some column CN to be known not nullable and no other constraint causes CN to be known not nullable, then the nullability characteristic of the column descriptor of CN is changed to known not nullable.

NOTE 298 — The nullability characteristic of a column is defined in [Subclause 4.13](#), “Columns, fields, and attributes”.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <add table constraint definition>.

11.21 <drop table constraint definition>

This Subclause is modified by *Subclause 9.9, “<drop table constraint definition>”, in ISO/IEC 9075-4.*

Function

Destroy a constraint on a table.

Format

```
<drop table constraint definition> ::=  
    DROP CONSTRAINT <constraint name> <drop behavior>
```

Syntax Rules

- 1) Let T be the table identified by the <table name> in the containing <alter table statement>.
- 2) The <constraint name> shall identify a table constraint TC of T .
- 3) If TC is a unique constraint and RC is a referential constraint whose referenced table is T and whose referenced columns are the unique columns of TC , then RC is said to be *dependent on TC* .
- 4) If QS is a <query specification> that contains an implicit or explicit <group by clause> and that contains a column reference to a column C in its <select list> that is not contained in an aggregated argument of a <set function specification>, and if G is the set of grouping columns of QS , and if the table constraint TC is needed to conclude that $G \rightarrow C$ is a known functional dependency in QS , then QS is said to be *dependent on TC* .
- 5) If V is a view that contains a <query specification> that is dependent on a table constraint TC , then V is said to be *dependent on TC* .
- 6) If R is an SQL routine whose <SQL routine body> contains a <query specification> that is dependent on a table constraint TC , then R is said to be *dependent on TC* .
- 7) If C is a constraint or assertion whose <search condition> contains a <query specification> that is dependent on a table constraint TC , then C is said to be *dependent on TC* .
- 8) If T is a trigger whose triggered action contains a <query specification> that is dependent on a table constraint TC , then T is said to be *dependent on TC* .
- 9) If T is a referenceable table with a derived self-referencing column, then:
 - a) TC shall not be a unique constraint whose unique columns correspond to the attributes in the list of attributes of the derived representation of the reference type whose referenced type is the structured type of T .
 - b) TC shall not be a unique constraint whose unique column is the self-referencing column of T .
- 10) If RESTRICT is specified, then:
 - a) No table constraint shall be dependent on TC .

11.21 <drop table constraint definition>

- b) The <constraint name> of *TC* shall not be generally contained in the SQL routine body of any routine descriptor.
- c) No view shall be dependent on *TC*.
- d) No SQL routine shall be dependent on *TC*.
- e) No constraint or assertion shall be dependent on *TC*.
- f) No trigger shall be dependent on *TC*.

NOTE 299 — If CASCADE is specified, then any such dependent object will be dropped by the effective execution of the General Rules of this Subclause.

Access Rules

None.

General Rules

- 1) Let *TCN2* be the <constraint name> of any table constraint that is dependent on *TC* and let *T2* be the <table name> of the table descriptor that includes *TCN2*. The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE T2 DROP
CONSTRAINT TCN2 CASCADE
```

- 2) 04 Let *R* be any SQL-invoked routine whose routine descriptor contains the <constraint name> of *TC* in the SQL routine body. Let *SN* be the <specific name> of *R*. The following <drop routine statement> is effectively executed without further Access Rule checking:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

- 3) Let *VN* be the table name of any view *V* that is dependent on *TC*. The following <drop view statement> is effectively executed for every *V*:

```
DROP VIEW VN CASCADE
```

- 4) Let *SN* be the specific name of any SQL routine *R* that is dependent on *TC*. The following <drop routine statement> is effectively executed for every *SR*:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

- 5) Let *CN* be the constraint name of any constraint *C* that is dependent on *TC*. Let *TN* be the name of the table constrained by *C*. The following <alter table statement> is effectively executed for every *C*:

```
ALTER TABLE TN DROP
CONSTRAINT CN CASCADE
```

- 6) Let *AN* be the assertion name of any assertion *A* that is dependent on *TC*. The following <drop assertion statement> is effectively executed for every *A*:

```
DROP ASSERTION AN CASCADE
```

11.21 <drop table constraint definition>

- 7) Let *TN* be the trigger name of any trigger *T* that is dependent on *TC*. The following <drop trigger statement> is effectively executed for every *T*:

```
DROP TRIGGER T CASCADE
```

- 8) The descriptor of *TC* is removed from the descriptor of *T*.
- 9) If *TC* causes some column *CN* to be known not nullable and no other constraint causes *CN* to be known not nullable, then the nullability characteristic of the column descriptor of *CN* is changed to possibly nullable.

NOTE 300 — The nullability characteristic of a column is defined in [Subclause 4.13](#), “Columns, fields, and attributes”.

- 10) The descriptor of *TC* is destroyed.

Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop table constraint definition>.

11.22 <drop table statement>

This Subclause is modified by *Subclause 9.10*, “<drop table statement>”, in ISO/IEC 9075-4.

Function

Destroy a table.

Format

```
<drop table statement> ::=
  DROP TABLE <table name> <drop behavior>
```

Syntax Rules

- 1) Let T be the table identified by the <table name> and let TN be that <table name>.
- 2) The schema identified by the explicit or implicit schema name of the <table name> shall include the descriptor of T .
- 3) T shall be a base table.
- 4) T shall not be a declared local temporary table.
- 5) An *impacted dereference operation* is any of the following:
 - a) A <dereference operation> DO , where T is the scope of the reference type of the <reference value expression> immediately contained in DO .
 - b) A <method reference> MR , where T is the scope of the reference type of the <value expression primary> immediately contained in MR .
 - c) A <reference resolution> RR , where T is the scope of the reference type of the <reference value expression> immediately contained in RR .
- 6) If RESTRICT is specified, then T shall not have any proper subtables.
- 7) 04 If RESTRICT is specified, then T shall not be referenced and no impacted dereference operation shall be contained in any of the following:
 - a) The <query expression> of any view descriptor.
 - b) The <search condition> of any constraint descriptor that is not a table check constraint descriptor included in the base table descriptor of T .
 - c) The <search condition> of any assertion descriptor.
 - d) The table descriptor of the referenced table of any referential constraint descriptor of any table other than T .
 - e) The SQL routine body of any routine descriptor.
 - f) The <triggered action> of any trigger descriptor.

NOTE 301 — If CASCADE is specified, then such objects will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.

- 8) If RESTRICT is specified and *T* is a referenceable table, then *TN* shall not be the scope included in a reference type descriptor generally included in any of the following:
- The attribute descriptor of an attribute of a user-defined type.
 - The column descriptor of a column of a table other than *T*.
 - The descriptor of an SQL parameter or the result type included in the routine descriptor of any <SQL-invoked routine>.
 - The descriptor of an SQL parameter or the result type included in a method specification descriptor included in the user-defined type descriptor of any user-defined type.
 - The descriptor of any user-defined cast.

NOTE 302 — A descriptor that “generally includes” another descriptor is defined in Subclause 6.3.4, “Descriptors”, in [ISO9075-1].

- 9) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of the table identified by *TN*.

Access Rules

- The enabled authorization identifiers shall include *A*.

General Rules

- Let *STN* be the <table name> of any direct subtable of *T*. The following <drop table statement> is effectively executed without further Access Rule checking:

```
DROP TABLE STN CASCADE
```

- For every proper supertable of *T*, every superrow of every row of *T* is effectively deleted at the end of the SQL-statement, prior to the checking of any integrity constraints.

NOTE 303 — This deletion creates neither a new trigger execution context nor the definition of a new state change in the current trigger execution context.

- The following <revoke statement> is effectively executed with a current authorization identifier of “_SYSTEM” and without further Access Rule checking:

```
REVOKE ALL PRIVILEGES ON TN FROM A CASCADE
```

- If *T* is a referenceable table, then:
 - For every reference type descriptor *RTD* that includes a scope of *TN* and is generally included in any of the following:
 - The attribute descriptor of an attribute of a user-defined type.
 - The column descriptor of a column of a table other than *T*.

11.22 <drop table statement>

- iii) The descriptor of an SQL parameter or the result type included in the routine descriptor of any SQL-invoked routine.
- iv) The descriptor of an SQL parameter or the result type in a method specification descriptor included in the user-defined type descriptor of any user-defined type.
- v) The descriptor of any user-defined cast.

the scope of *RTD* is made empty.

- b) Let *SOD* be the descriptor of a schema object dependent on the table descriptor of *T*.

Case:

- i) If *SOD* is a view descriptor, then let *SON* be the name of the view included in *SOD*. The following <drop view statement> is effectively executed without further Access Rule checking:

```
DROP VIEW SON CASCADE
```

- ii) If *SOD* is an assertion descriptor, then let *SON* be the name of the assertion included in *SOD*. The following <drop assertion statement> is effectively executed without further Access Rule checking:

```
DROP ASSERTION SON CASCADE
```

- iii) If *SOD* is a table constraint descriptor, then let *SON* be the name of the constraint included in *SOD*. Let *CTN* be the <table name> included in the table descriptor that includes *SOD*. The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE CTN DROP CONSTRAINT SON CASCADE
```

- iv) If *SOD* is a routine descriptor, then let *SON* be the specific name included in *SOD*. The following <drop routine statement> is effectively executed without further Access Rule checking:

```
DROP SPECIFIC ROUTINE SON CASCADE
```

- v) If *SOD* is a trigger descriptor, then let *SON* be the trigger name included in *SOD*. The following <drop trigger statement> is effectively executed without further Access Rule checking:

```
DROP TRIGGER SON CASCADE
```

NOTE 304 — A descriptor that “depends on” another descriptor is defined in [Subclause 6.3.4, “Descriptors”](#), in [ISO9075-1].

- 5) For each direct supertable *DST* of *T*, the table name of *T* is removed from the list of table names of direct subtables of *DST* that is included in the table descriptor of *DST*.
- 6) The descriptor of *T* is destroyed.

Conformance Rules

- 1) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop table statement> that contains <drop behavior> that contains CASCADE.

11.23 <view definition>

This Subclause is modified by Subclause 9.11, “<view definition>”, in ISO/IEC 9075-4.

This Subclause is modified by Subclause 9.3, “<view definition>”, in ISO/IEC 9075-13.

This Subclause is modified by Subclause 12.3, “<view definition>”, in ISO/IEC 9075-14.

Function

Define a viewed table.

Format

```

<view definition> ::=
    CREATE [ RECURSIVE ] VIEW <table name> <view specification>
        AS <query expression> [ WITH [ <levels clause> ] CHECK OPTION ]

<view specification> ::=
    <regular view specification>
    | <referenceable view specification>

<regular view specification> ::=
    [ <left paren> <view column list> <right paren> ]

<referenceable view specification> ::=
    OF <path-resolved user-defined type name> [ <subview clause> ]
        [ <view element list> ]

<subview clause> ::=
    UNDER <table name>

<view element list> ::=
    <left paren> <view element> [ { <comma> <view element> }... ] <right paren>

<view element> ::=
    <self-referencing column specification>
    | <view column option>

<view column option> ::=
    <column name> WITH OPTIONS <scope clause>

<levels clause> ::=
    CASCADED
    | LOCAL

<view column list> ::=
    <column name list>

```

Syntax Rules

- 1) 04 The <query expression> shall not contain a <host parameter specification>, an <SQL parameter reference>, a <dynamic parameter specification>, or an <embedded variable specification>.

- 2) If a <view definition> is contained in a <schema definition> and the <table name> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the containing <schema definition>.
- 3) The schema identified by the explicit or implicit schema name of the <table name> shall not include a table descriptor whose table name is <table name>.
- 4) No <table reference> generally contained in the <query expression> shall identify any declared local temporary table.
- 5) If a <table reference> generally contained in the <query expression> identifies the viewed table *VT* defined by <view definition> *VD*, then *VD* and *VT* are said to be *recursive*.
- 6) If *VD* is recursive, then:
 - a) <view column list> shall be specified.
 - b) **RECURSIVE** shall be specified.
 - c) **CHECK OPTION** shall not be specified.
 - d) <referenceable view specification> shall not be specified.
 - e) *VD* is equivalent to

```
CREATE VIEW <table name> AS
  WITH RECURSIVE <table name> (<view column list>)
  AS (<query expression>)
  SELECT <view column list> FROM <table name>
```

- 7) The viewed table is *updatable* if the <query expression> is updatable.
- 8) The viewed table is *simply updatable* if the <query expression> is simply updatable.
- 9) The viewed table is *effectively updatable* if it is simply updatable, or if the SQL implementation supports Feature T111, “Updatable joins, unions, and columns”, and the viewed table is updatable.
- 10) The viewed table is *insertable-into* if the <query expression> is insertable-into.
- 11) If any two columns in the table specified by the <query expression> have equivalent <column name>s, or if any column of that table has an implementation-dependent name, then a <view column list> shall be specified.
- 12) Equivalent <column name>s shall not be specified more than once in the <view column list>.
- 13) The number of <column name>s in the <view column list> shall be the same as the degree of the table specified by the <query expression>.
- 14) Every column in the table specified by <query expression> whose declared type is a character string type shall have a declared type collation.
- 15) If **WITH CHECK OPTION** is specified, then the viewed table shall be updatable.
- 16) If **WITH CHECK OPTION** is specified and <levels clause> is not specified, then a <levels clause> of **CASCADED** is implicit.

- 17) If WITH LOCAL CHECK OPTION is specified, then the <query expression> shall not generally contain a <query expression> *QE* or a <query specification> *QS* that is *possibly non-deterministic* unless *QE* or *QS* is generally contained in a viewed table that is a leaf underlying table of the <query expression>.

If WITH CASCADED CHECK OPTION is specified, then the <query expression> shall not generally contain a <query expression> or <query specification> that is *possibly non-deterministic*.

- 18) Let *V* be the view defined by the <view definition>. The underlying columns of every *i*-th column of *V* are the underlying columns of the *i*-th column of the <query expression> and the underlying columns of *V* are the underlying columns of the <query expression>.

- 19) <subview clause>, if present, identifies the *direct superview* *SV* of *V*; *V* is said to be a *direct subview* of *SV*. View *V1* is a *superview* of view *V2* if and only if one of the following is true:

- a) *V1* and *V2* are the same view.
- b) *V1* is a direct superview of *V2*.
- c) There exists a view *V3* such that *V1* is a direct superview of *V3* and *V3* is a superview of *V2*.

If *V1* is a superview of *V2*, then *V2* is said to be a *subview* of *V1*.

If *V1* is a superview of *V2* and *V1* and *V2* are not the same view, then *V2* is a *proper subview* of *V1* and *V1* is a *proper superview* of *V2*.

If *V2* is a direct subview of *V1*, then *V2* is a direct subtable of *V1*.

NOTE 305 — It follows that the subviews of the supervIEWS of *V* together constitute the subtable family of *V*, every implication of which applies.

- 20) If <referenceable view specification> is specified, then:

- a) *V* is a *referenceable view*.
- b) RECURSIVE shall not be specified.
- c) 13 The <user-defined type name> simply contained in <path-resolved user-defined type name> shall identify a structured type *ST*.
- d) The subtable family of *V* shall not include a member, other than *V* itself, whose associated structured type is *ST*.
- e) If <subview clause> is not specified, then <self-referencing column specification> shall be specified.
- f) Let *QE* be the <query expression>.
- g) Let *n* be the number of attributes of *ST*. Let *A_i*, 1 (one) ≤ *i* ≤ *n* be the attributes of *ST*.
- h) Let *RT* be the row type of *QE*.
- i) If <self-referencing column specification> is specified, then:
 - i) Exactly one <self-referencing column specification> shall be specified.
 - ii) <subview clause> shall not be specified.
 - iii) SYSTEM GENERATED shall not be specified.
 - iv) Let *RST* be the reference type REF(*ST*).

Case:

- 1) If USER GENERATED is specified, then:
 - A) *RST* shall have a user-defined representation.
 - B) Let m be 1 (one).
- 2) If DERIVED is specified, then:
 - A) *RST* shall have a derived representation.
 - B) Let m be 0 (zero).
- j) If <subview clause> is specified, then:
 - i) The <table name> contained in the <subview clause> shall identify a referenceable table *SV* that is a view.
 - ii) *ST* shall be a direct subtype of the structured type of the direct supertable of *V*.
 - iii) The SQL-schema identified by the explicit or implicit <schema name> of the <table name> of *V* shall include the descriptor of *SV*.
 - iv) Let *MSV* be the maximum superview of the subtable family of *V*. Let *RMSV* be the reference type REF(*MSV*).

Case:

- 1) If *RMSV* has a user-defined representation, then let m be 1 (one).
- 2) Otherwise, *RMSV* has a derived representation. Let m be 0 (zero).
- k) The degree of *RT* shall be $n+m$.
- l) Let F_i , $1 \text{ (one)} \leq i \leq n$, be the fields of *RT*.
- m) For i varying from 1 (one) to n :
 - i) The declared type $DDTF_{i+m}$ of F_{i+m} shall be compatible with the declared type $DDTA_i$ of A_i .
 - ii) The Syntax Rules of Subclause 9.18, “Data type identity”, are applied with $DDTF_{i+m}$ and $DDTA_i$.
- n) *QE* shall consist of a single <query specification> *QS*.
- o) The <from clause> of *QS* shall simply contain a single <table reference> *TR*.
- p) *TR* shall immediately contain a <table or query name>. Let *TQN* be the table identified by the <table or query name>. *TQN* is the *basis table* of *V*.
- q) If *TQN* is a referenceable base table or a referenceable view, then *TR* shall simply contain ONLY.
- r) *QS* shall not simply contain a <group by clause> or a <having clause>.
- s) If <self-referencing column specification> is specified, then

Case:

- i) If *RST* has a user-defined representation, then:

- 1) TQN shall have a candidate key consisting of a single column RC .
- 2) Let SS be the first <select sublist> in the <select list> of QS .
- 3) SS shall consist of a single <cast specification> CS whose leaf column is RC .
NOTE 306 — “Leaf column of a <cast specification>” is defined in Subclause 6.12, “<cast specification>”.
- 4) The declared type of F_1 shall be $REF(ST)$.
- ii) Otherwise, RST has a derived representation.
 - 1) Let C_i , $1 \text{ (one)} \leq i \leq n$, be the columns of V that correspond to the attributes of the derived representation of RST .
 - 2) TQN shall have a candidate key consisting of some subset of the underlying columns of C_i , $1 \text{ (one)} \leq i \leq n$.
- t) If <subview clause> is specified, then TQN shall be a proper subtable or proper subview of the basis table of SV .
- u) If <view element list> $TELI$ is specified, then:
 - i) Let r be the number of <view column option>s. For every <view column option> VCO_j , $1 \text{ (one)} \leq j \leq r$, <column name> shall be equivalent to the <attribute name> of some attribute of ST .
 - ii) Distinct <view column option>s contained in $TELI$ shall specify distinct <column name>s.
 - iii) Let CN_j , $1 \text{ (one)} \leq j \leq r$, be the <column name> contained in VCO_j and let SCL_j be the <scope clause> contained in VCO_j .
 - 1) CN_j shall be equivalent to some <attribute name> of ST , whose declared type is some reference type $CORT_j$.
 - 2) The <table name> contained in SCL_j shall identify a referenceable table SRT .
 - 3) SRT shall be based on the referenced type of $CORT_j$.

NOTE 307 — The notion of one data type being based on another data type is defined in Subclause 4.1, “Data types”.

- 21) Let the *originally-defined columns* of V be the columns of the table defined by QE .
- 22) A column of V is called an *updatable column* of V if its underlying column is updatable.
- 23) If the <view definition> is contained in a <schema definition>, then let A be the explicit or implicit <authorization identifier> of the <schema definition>; otherwise, let A be the <authorization identifier> that owns the schema identified by the explicit or implicit <schema name> of the <table name>.

Access Rules

- 1) If a <view definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of the <table name>.
- 2) If <referenceable view specification> is specified, then the applicable privileges for A shall include USAGE on ST .

3) If <subview clause> is specified, then

Case:

- a) If <view definition> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the schema shall include UNDER for SV.
- b) Otherwise, the current privileges shall include UNDER for SV.

General Rules

1) A view descriptor *VD* is created that describes *V*. *VD* includes:

- a) The <table name> *TN*.
- b) *QE*, as both the <query expression> of the descriptor and the original <query expression> of the descriptor.
- c) Case:
 - i) If <regular view specification> is specified, then the column descriptors taken from the table specified by the <query expression>.

Case:

- 1) If a <view column list> is specified, then the <column name> of the *i*-th column of the view is the *i*-th <column name> in that <view column list>.
- 2) Otherwise, the <column name>s of the view are the <column name>s of the table specified by the <query expression>.
- ii) Otherwise:
 - 1) A column descriptor in which:
 - A) The name of the column is <self-referencing column name>.
 - B) The data type descriptor is that generated by the <data type> “REF(*ST*) SCOPE(*TN*)”.
 - C) The nullability characteristic is *known not nullable*.
 - D) The ordinal position is 1 (one).
 - E) The column is indicated to be self-referencing.
 - 2) The column descriptor *ODCD* of each originally-defined column *ODC* of *V* in which:
 - A) The <column name> included in *ODCD* is replaced by the <attribute name> of its corresponding attribute of *ST*.
 - B) If the declared type of the column is a reference type and some *VCO_i* contains the <attribute name> of *ST* that corresponds to the column, then the (possibly empty) scope contained in the reference type descriptor immediately included in the column descriptor is replaced by *SCO_i*.

- 3) If DERIVED is specified, then an indication that the self-referencing column is a derived self-referencing column.
- 4) If USER GENERATED is specified, then an indication that the self-referencing column is a user-generated self-referencing column.
- d) In each column descriptor, an indication that the column is updatable if *V* is effectively updatable and the corresponding column of *QE* is updatable.
- e) An indication as to whether WITH CHECK OPTION was omitted, specified with LOCAL, or specified with CASCADED.
- f) An indication that *V* is not trigger insertable-into.
- g) An indication that *V* is not trigger updatable.
- h) An indication that *V* is not trigger deletable.
- 2) Let *VN* be the <table name>. Let *QE* be the <query expression> included in the view descriptor *VD* of the view identified by *VN*. Let *OQE* be the original <query expression> included in *VD*. If a <view column list> is specified, then let *VCL* be the <view column list> preceded by a <left paren> and followed by a <right paren>; otherwise, let *VCL* be the zero-length string.

Case:

- a) If *VN* is immediately contained in some SQL-schema statement, then *VN* identifies the view descriptor *VD*.
- b) If *VN* is immediately contained in a <table reference> that specifies ONLY, then *VN* references the same table as the <table reference>:

(*OQE*) AS *VN* *VCL*

- c) Otherwise, *VN* references the same table as the <table reference>:

(*QE*) AS *VN* *VCL*

- 3) For *i* ranging from 1 (one) to the number of distinct leaf underlying tables of the <query expression> *QE* of *V*, let *RT_i* be the <table name>s of those tables. For every column *CV* of *V*:
 - a) Let *CRT_{i,j}*, for *j* ranging from 1 (one) to the number of columns of *RT_i* that are underlying columns of *CV*, be the <column name>s of those columns.
 - b) A set of privilege descriptors with the grantor for each set to the special grantor value “_SYSTEM” is created as follows:
 - i) For every column *CV* of *V*, a privilege descriptor is created that defines the privilege SELECT(*CV*) on *V* to *A*. That privilege is grantable if and only if all the following are true:
 - 1) The applicable privileges for *A* include grantable SELECT privileges on all of the columns *CRT_{i,j}*.
 - 2) The applicable privileges for *A* include grantable EXECUTE privileges on all SQL-invoked routines that are subject routines of <routine invocation>s contained in *QE*.

- 3) The applicable privileges for *A* include grantable SELECT privilege on every table *T_l* and every method *M* such that there is a <method reference> *MR* contained in *QE* such that *T_l* is in the scope of the <value expression primary> of *MR* and *M* is the method identified by the <method name> of *MR*.
 - 4) The applicable privileges for *A* include grantable SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of every <reference resolution> that is contained in *QE*.
- ii) For every column *CV* of *V*, if the applicable privileges for *A* include REFERENCES(*CRT_{i,j}*) for all *i* and for all *j*, and the applicable privileges for *A* include REFERENCES on some column of *RT_i* for all *i*, then a privilege descriptor is created that defines the privilege REFERENCES(*CV*) on *V* to *A*. That privilege is grantable if and only if all the following are true:
- 1) The applicable privileges for *A* include grantable REFERENCES privileges on all of the columns *CRT_{i,j}*.
 - 2) The applicable privileges for *A* include grantable EXECUTE privileges on all SQL-invoked routines that are subject routines of <routine invocation>s contained in *QE*.
 - 3) The applicable privileges for *A* include grantable SELECT privilege on every table *T_l* and every method *M* such that there is a <method reference> *MR* contained in *QE* such that *T_l* is in the scope of the <value expression primary> of *MR* and *M* is the method identified by the <method name> of *MR*.
 - 4) The applicable privileges for *A* include grantable SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of every <reference resolution> that is contained in *QE*.
- 4) A privilege descriptor is created that defines the privilege SELECT on *V* to *A*. That privilege is grantable if and only if the applicable privileges for *A* include grantable SELECT privilege on every column of *V*. The grantor of that privilege descriptor is set to the special grantor value “_SYSTEM”.
 - 5) A privilege descriptor is created that defines the privilege TRIGGER on *V* to *A*. That privilege is not grantable. The grantor of that privilege descriptor is set to the special grantor value “_SYSTEM”.
 - 6) If the applicable privileges for *A* include REFERENCES privilege on every column of *V*, then a privilege descriptor is created that defines the privilege REFERENCES on *V* to *A*. That privilege is grantable if and only if the applicable privileges for *A* include grantable REFERENCES privilege on every column of *V*. The grantor of that privilege descriptor is set to the special grantor value “_SYSTEM”.
 - 7) If *V* is effectively updatable, then the General Rules of Subclause 10.11, “Determination of view privileges”, are applied with *V* as *VIEW*, yielding a set of privilege descriptors returned from *DESCRIPTOR SET*.
 - 8) If *V* is a referenceable view, then a set of privilege descriptors with the grantor for each set to the special grantor value “_SYSTEM” are created as follows:
 - a) A privilege descriptor is created that defines the SELECT privilege WITH HIERARCHY OPTION on *V* to *A*. That privilege is grantable.
 - b) For every method *M* of the structured type identified by <path-resolved user-defined type name>, a privilege descriptor is created that defines the privilege SELECT(*M*) on *V* to *A*. That privilege is grantable.
 - c) Case:

- i) If <subview clause> is not specified, then a privilege descriptor is created that defines the UNDER privilege on *V* to *A*. That privilege is grantable.
 - ii) Otherwise, a privilege descriptor is created that defines the UNDER privilege on *V* to *A*. That privilege is grantable if and only if the applicable privileges for *A* include grantable UNDER privilege on the direct supertable of *V*.
- 9) If <subview clause> is specified, then let *ST* be the set of supertables of *V*. Let *PDS* be the set of privilege descriptors that define SELECT WITH HIERARCHY OPTION privilege on a table in *ST*.
- 10) For every privilege descriptor in *PDS*, with grantee *G* and grantor *A*,
- Case:
- a) If the privilege is grantable, then let *WGO* be “WITH GRANT OPTION”.
 - b) Otherwise, let *WGO* be a zero-length string.

The following <grant statement> is effectively executed without further Access Rule checking:

```
GRANT SELECT ON V TO G WGO FROM A
```

- 11) If <subview clause> is specified, then let *SVQE* be the <query expression> included in the view descriptor of *SV*.
- a) The <query expression> included in the descriptor of *SV* is replaced by the following <query expression>:
- ```
(SVQE) UNION ALL CORRESPONDING SELECT * FROM TN
```
- b) The General Rules of this subclause are reevaluated for *SV* in the light of the new <query expression> in its descriptor.
- 12) If the character representation of the <query expression> cannot be represented in the Information Schema without truncation, then a completion condition is raised: *warning — query expression too long for information schema*.

NOTE 308 — The Information Schema is defined in [ISO9075-11].

## Conformance Rules

- 1) Without Feature T131, “Recursive query”, conforming SQL language shall not contain a <view definition> that immediately contains RECURSIVE.
- 2) Without Feature F751, “View CHECK enhancements”, conforming SQL language shall not contain a <levels clause>.
- 3) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <referenceable view specification>.
- 4) Without Feature F751, “View CHECK enhancements”, conforming SQL language shall not contain <view definition> that contains a <query expression> that contains a <query expression> and contains WITH CHECK OPTION.
- 5) Without Feature T111, “Updatable joins, unions and columns”, in conforming SQL language, if WITH CHECK OPTION is specified, then the viewed table shall be simply updatable.

- 6) Without Feature F852, “Top-level <order by clause> in views”, in conforming SQL language, a <query expression> immediately contained in a <view definition> shall not immediately contain an <order by clause>.
- 7) Without Feature F859, “Top-level <fetch first clause> in views”, in conforming SQL language, a <query expression> immediately contained in a <view definition> shall not immediately contain a <fetch first clause>.

## 11.24 <drop view statement>

This Subclause is modified by *Subclause 9.12, “<drop view statement>”, in ISO/IEC 9075-4.*

### Function

Destroy a view.

### Format

```
<drop view statement> ::=
 DROP VIEW <table name> <drop behavior>
```

### Syntax Rules

- 1) Let *V* be the table identified by the <table name> and let *VN* be that <table name>. The schema identified by the explicit or implicit schema name of *VN* shall include the descriptor of *V*.
- 2) *V* shall be a viewed table.
- 3) An *impacted dereference operation* is any of the following:
  - a) A <dereference operation> *DO*, where *V* is the scope of the reference type of the <reference value expression> immediately contained in *DO*.
  - b) A <method reference> *MR*, where *V* is the scope of the reference type of the <value expression primary> immediately contained in *MR*.
  - c) A <reference resolution> *RR*, where *V* is the scope of the reference type of the <reference value expression> immediately contained in *RR*.
- 4) If RESTRICT is specified, then *V* shall not have any proper subtables.
- 5) 04 If RESTRICT is specified, then *V* shall not be referenced and no impacted dereference operation shall be contained in any of the following:
  - a) The <query expression> of the view descriptor of any view other than *V*.
  - b) The <search condition> of any constraint descriptor or assertion descriptor.
  - c) The <triggered action> of any trigger descriptor.
  - d) The SQL routine body of any routine descriptor.

NOTE 309 — If CASCADE is specified, then any such dependent object will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.
- 6) If RESTRICT is specified and *V* is a referenceable view, then *VN* shall not be the scope included in a reference type descriptor generally included in any of the following:
  - a) The attribute descriptor of an attribute of a user-defined type.
  - b) The column descriptor of a column of a table other than *V*.

- c) The descriptor of an SQL parameter or the result type included in the routine descriptor of any <SQL-invoked routine>.
- d) The descriptor of an SQL parameter or the result type included in a method specification descriptor included in the user-defined type descriptor of any user-defined type.
- e) The descriptor of any user-defined cast.

NOTE 310 — A descriptor that “generally includes” another descriptor is defined in Subclause 6.3.4, “Descriptors”, in [ISO9075-1].

- 7) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of the table identified by *VN*.

## Access Rules

- 1) The enabled authorization identifier shall include *A*.

## General Rules

- 1) Let *SVN* be the <table name> of any direct subview of *V*. The following <drop view statement> is effectively executed without further Access Rule checking:

```
DROP VIEW SVN CASCADE
```

- 2) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE ALL PRIVILEGES ON VN FROM A CASCADE
```

- 3) If *V* is a referenceable view, then:

- a) For every reference type descriptor *RTD* that includes a scope of *VN* and is generally included in any of the following:
  - i) The attribute descriptor of an attribute of a user-defined type.
  - ii) The column descriptor of a column of a table other than *V*.
  - iii) The descriptor of an SQL parameter or the result type included in the routine descriptor of any SQL-invoked routine.
  - iv) The descriptor of an SQL parameter or the result type included in a method specification descriptor included in the user-defined type descriptor of any user-defined type.
  - v) The descriptor of any user-defined cast.

the scope of *RTD* is made empty.

- b) Let *SOD* be the descriptor of a schema object dependent on the view descriptor of *V*.

Case:

- i) If *SOD* is a view descriptor, then let *SON* be the name of the view included in *SOD*. The following <drop view statement> is effectively executed without further Access Rule checking:

DROP VIEW *SON* CASCADE

- ii) If *SOD* is an assertion descriptor, then let *SON* be the name of the assertion included in *SOD*. The following <drop assertion statement> is effectively executed without further Access Rule checking:

DROP ASSERTION *SON* CASCADE

- iii) If *SOD* is a table constraint descriptor, then let *SON* be the name of the constraint included in *SOD*. Let *CTN* be the <table name> included in the table descriptor that includes *SOD*. The following <alter table statement> is effectively executed without further Access Rule checking:

ALTER TABLE *CTN* DROP CONSTRAINT *SON* CASCADE

- iv) If *SOD* is a routine descriptor, then let *SON* be the specific name included in *SOD*. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE *SON* CASCADE

- v) If *SOD* is a trigger descriptor, then let *SON* be the trigger name included in *SOD*. The following <drop trigger statement> is effectively executed without further Access Rule checking:

DROP TRIGGER *SON*

NOTE 311 — A descriptor that “depends on” another descriptor is defined in Subclause 6.3.4, “Descriptors”, in [ISO9075-1].

- 4) For each direct supertable *DST* of *V*, the table name of *V* is removed from the list of table names of direct subtables of *DST* that is included in the table descriptor of *DST*.
- 5) The descriptor of *V* is destroyed.

## Conformance Rules

- 1) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop view statement> that contains a <drop behavior> that contains CASCADE.

## 11.25 <domain definition>

*This Subclause is modified by Subclause 11.7, “<domain definition>”, in ISO/IEC 9075-9.*

### Function

Define a domain.

### Format

```
<domain definition> ::=
 CREATE DOMAIN <domain name> [AS] <predefined type>
 [<default clause>]
 [<domain constraint>...]
 [<collate clause>]

<domain constraint> ::=
 [<constraint name definition>] <check constraint definition> [
 <constraint characteristics>]
```

### Syntax Rules

- 1) 09 If a <domain definition> is contained in a <schema definition>, and if the <domain name> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the containing <schema definition>.
- 2) If <constraint name definition> is specified and its <constraint name> contains a <schema name>, then that <schema name> shall be equivalent to the explicit or implicit <schema name> of the <domain name> of the domain identified by the containing <domain definition> or <alter domain statement>.
- 3) The schema identified by the explicit or implicit schema name of the <domain name> shall not include a domain descriptor whose domain name is equivalent to <domain name> nor a user-defined type descriptor whose user-defined type name is equivalent to <domain name>.
- 4) If <predefined type> specifies a <character string type> and does not specify <character set specification>, then the character set name of the default character set of the schema identified by the implicit or explicit <schema name> of <domain name> is implicit.
- 5) <collate clause> shall not be both specified in <predefined type> and immediately contained in <domain definition>. If <collate clause> is immediately contained in <domain definition>, then it is equivalent to specifying an equivalent <collate clause> in <predefined type>.
- 6) Let *D1* be some domain. *D1* is *in usage* by a domain constraint *DC* if and only if the <search condition> of *DC* generally contains the <domain name> either of *D1* or of some domain *D2* such that *D1* is in usage by some domain constraint of *D2*. No domain shall be in usage by any of its own constraints.
- 7) If <collate clause> is specified, then <predefined type> shall be a character string type.
- 8) For every <domain constraint> specified:
  - a) If <constraint characteristics> is not specified, then INITIALLY IMMEDIATE NOT DEFERRABLE is implicit.



- b) If <constraint name definition> is not specified, then a <constraint name definition> that contains an implementation-dependent <constraint name> is implicit. The assigned <constraint name> shall obey the Syntax Rules of an explicit <constraint name>.
- c) Let *S* be the schema identified by the explicit or implicit <schema name> of the <constraint name> contained in <domain constraint>. *S* shall not include a constraint descriptor whose constraint name is <constraint name>.

## Access Rules

- 1) If a <domain definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of the <domain name>.

## General Rules

- 1) A <domain definition> defines a domain.
- 2) A data type descriptor is created that describes the declared type of the domain being created.
- 3) A domain descriptor is created that describes the domain being created. The domain descriptor contains the name of the domain, the data type descriptor of the declared type, the value of the <default clause> if the <domain definition> immediately contains <default clause>, and a domain constraint descriptor for every immediately contained <domain constraint>.
- 4) A privilege descriptor is created that defines the USAGE privilege on this domain to the <authorization identifier> *A* of the schema or SQL-client module in which the <domain definition> appears. This privilege is grantable if and only if the applicable privileges for *A* include a grantable REFERENCES privilege for each column reference included in the domain descriptor and a grantable USAGE privilege for each <domain name>, <collation name>, <character set name>, and <transliteration name> contained in the <search condition> of any domain constraint descriptor included in the domain descriptor. The grantor of the privilege descriptor is set to the special grantor value “\_SYSTEM”.
- 5) A domain constraint descriptor is created that describes the domain constraint being defined. The domain constraint descriptor includes:
  - a) The <constraint name> contained in the explicit or implicit <constraint name definition>.
  - b) An indication of whether the constraint is deferrable or not deferrable.
  - c) An indication of whether the initial constraint mode of the constraint is deferred or immediate.
  - d) The <search condition> contained in the <domain definition> as the template <search condition>.
  - e) The applicable <search condition>:

( 1=1 )

## Conformance Rules

- 1) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <domain definition>.

- 2) Without Feature F692, “Extended collation support”, conforming SQL language shall not contain a <domain definition> that immediately contains a <collate clause>.

## 11.26 <alter domain statement>

### Function

Change a domain and its definition.

### Format

```
<alter domain statement> ::=
 ALTER DOMAIN <domain name> <alter domain action>

<alter domain action> ::=
 <set domain default clause>
 | <drop domain default clause>
 | <add domain constraint definition>
 | <drop domain constraint definition>
```

### Syntax Rules

- 1) Let  $D$  be the domain identified by <domain name>. The schema identified by the explicit or implicit schema name of the <domain name> shall include the descriptor of  $D$ .

### Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of <domain name>.

### General Rules

- 1) The domain descriptor of  $D$  is modified as specified by <alter domain action>.

NOTE 312 — The changed domain descriptor of  $D$  is applicable to every column that is dependent on  $D$ .

### Conformance Rules

- 1) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain an <alter domain statement>.

## 11.27 <set domain default clause>

### Function

Set the default value in a domain.

### Format

```
<set domain default clause> ::=
 SET <default clause>
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let  $D$  be the domain identified by the <domain name> in the containing <alter domain statement>.
- 2) The default value specified by the <default clause> is placed in the domain descriptor of  $D$ .

### Conformance Rules

- 1) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain a <set domain default clause>.

## 11.28 <drop domain default clause>

### Function

Remove the default clause of a domain.

### Format

```
<drop domain default clause> ::=
 DROP DEFAULT
```

### Syntax Rules

- 1) Let  $D$  be the domain identified by the <domain name> in the containing <alter domain statement>.
- 2) The descriptor of  $D$  shall contain a default value.

### Access Rules

*None.*

### General Rules

- 1) Let  $C$  be the set of columns whose column descriptors contain the <domain name> that identifies  $D$ .
- 2) For every column belonging to  $C$ , if the column descriptor does not already contain a default value, then the default value from the domain descriptor of  $D$  is placed in that column descriptor.
- 3) The default value is removed from the domain descriptor of  $D$ .

### Conformance Rules

- 1) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain a <drop domain default clause>.

## 11.29 <add domain constraint definition>

### Function

Add a constraint to a domain.

### Format

```
<add domain constraint definition> ::=
 ADD <domain constraint>
```

### Syntax Rules

- 1) Let *D* be the domain identified by the <domain name> in the <alter domain statement>.
- 2) Let *D1* be some domain. *D1* is *in usage* by a domain constraint *DC* if and only if the <search condition> of *DC* generally contains the <domain name> either of *D1* or of some domain *D2* such that *D1* is in usage by some domain constraint of *D2*. No domain shall be in usage by any of its own constraints.

### Access Rules

*None.*

### General Rules

- 1) The constraint descriptor of the <domain constraint> is added to the domain descriptor of *D*.
- 2) If *DC* causes some column *CN* to be known not nullable and no other constraint causes *CN* to be known not nullable, then the nullability characteristic of the column descriptor of *CN* is changed to known not nullable.

NOTE 313 — The nullability characteristic of a column is defined in Subclause 4.13, “Columns, fields, and attributes”.

### Conformance Rules

- 1) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain an <add domain constraint definition>.

## 11.30 <drop domain constraint definition>

### Function

Destroy a constraint on a domain.

### Format

```
<drop domain constraint definition> ::=
 DROP CONSTRAINT <constraint name>
```

### Syntax Rules

- 1) Let *D* be the domain identified by the <domain name> *DN* in the containing <alter domain statement>.
- 2) Let *CD* be any column descriptor that includes *DN*, let *T* be the table described by the table descriptor that includes *CD*, and let *TN* be the <table name> of *T*.
- 3) Let *DC* be the descriptor of the constraint identified by <constraint name>.
- 4) *DC* shall be included in the domain descriptor of *D*.

### Access Rules

*None.*

### General Rules

- 1) The constraint descriptor *DC* is removed from the domain descriptor of *D*.
- 2) If *DC* causes some column *CN* to be known not nullable and no other constraint causes *CN* to be known not nullable, then the nullability characteristic of the column descriptor of *CN* is changed to possibly nullable.  
 NOTE 314 — The nullability characteristic of a column is defined in [Subclause 4.13](#), “Columns, fields, and attributes”.
- 3) The descriptor of *DC* is destroyed.

### Conformance Rules

- 1) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain a <drop domain constraint definition>.
- 2) Without Feature F491, “Constraint management”, conforming SQL language shall not contain a <drop domain constraint definition>.

## 11.31 <drop domain statement>

*This Subclause is modified by Subclause 9.13, “<drop domain statement>”, in ISO/IEC 9075-4.*

### Function

Destroy a domain.

### Format

```
<drop domain statement> ::=
 DROP DOMAIN <domain name> <drop behavior>
```

### Syntax Rules

- 1) Let  $D$  be the domain identified by <domain name> and let  $DN$  be that <domain name>. The schema identified by the explicit or implicit schema name of  $DN$  shall include the descriptor of  $D$ .
- 2) 04 If RESTRICT is specified, then  $D$  shall not be referenced in any of the following:
  - a) A column descriptor.
  - b) The <query expression> of any view descriptor.
  - c) The <search condition> of any constraint descriptor.
  - d) The SQL routine body of any routine descriptor.
- 3) Let  $UA$  be the <authorization identifier> that owns the schema identified by the <schema name> of the domain identified by  $DN$ .

### Access Rules

- 1) The enabled authorization identifiers shall include  $UA$ .

### General Rules

- 1) Let  $C$  be any column descriptor that includes  $DN$ , let  $T$  be the table described by the table descriptor that includes  $C$ , and let  $TN$  be the table name of  $T$ .  $C$  is modified as follows:
  - a)  $DN$  is removed from  $C$ . A copy of the data type descriptor of  $D$  is included in  $C$ .
  - b) If  $C$  does not include a <default clause> and the domain descriptor of  $D$  includes a <default clause>, then a copy of the <default clause> of  $D$  is included in  $C$ .
  - c) For every domain constraint descriptor  $DCD$  included in the domain descriptor of  $D$ , for every domain constraint usage descriptor  $DCU$  included in  $DCD$ :
    - i) Let  $SC$  be the applicable <search condition> included in  $DCU$ .



- ii) Let *TCD* be a <table constraint definition> consisting of a <constraint name definition> whose <constraint name> is implementation-dependent, whose <constraint characteristics> are the <constraint characteristics> of the domain constraint descriptor, and whose <table constraint> is:

```
CHECK (SC)
```

- iii) If the applicable privileges for *UA* include all of the privileges necessary for *UA* to successfully execute the <alter table statement>

```
ALTER TABLE TN ADD TCD
```

then the following <alter table statement> is effectively executed with a current authorization identifier of *UA*:

```
ALTER TABLE TN ADD TCD
```

- 2) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE USAGE ON DOMAIN DN
FROM UA CASCADE
```

- 3) The descriptor of *D* is destroyed.

## Conformance Rules

- 1) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <drop domain statement>.

## 11.32 <character set definition>

### Function

Define a character set.

### Format

```
<character set definition> ::=
 CREATE CHARACTER SET <character set name> [AS]
 <character set source> [<collate clause>]

<character set source> ::=
 GET <character set specification>
```

### Syntax Rules

- 1) If a <character set definition> is contained in a <schema definition> and if the <character set name> immediately contained in the <character set definition> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the <schema definition>.
- 2) The schema identified by the explicit or implicit schema name of the <character set name> shall not include a character set descriptor whose character set name is <character set name>.
- 3) The character set *CS* identified by the <character set specification> contained in <character set source> shall have associated with it a privilege descriptor that was effectively defined by the <grant statement>

```
GRANT USAGE ON CHARACTER
SET CSN TO PUBLIC
```

where *CSN* is a <character set name> that identifies *CS*.

- 4) If <collate clause> is specified, then the <collation name> contained in <collate clause> shall identify a collation descriptor *CD* included in the schema identified by the explicit or implicit <schema name> contained in the <collation name>. The collation shall be applicable to the character repertoire of the character set identified by <character set source>. The list of applicable character set names included in *CD* shall include one that identifies *CS*.
- 5) Let *A* be the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of <character set name>.

### Access Rules

- 1) If a <character set definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.
- 2) The applicable privileges for *A* shall include USAGE on the character set identified by the <character set specification>.

## General Rules

- 1) A <character set definition> defines a character set.
- 2) A character set descriptor is created for the defined character set.
- 3) The descriptor created for the character set being defined is identical to the descriptor for the character set identified by <character set specification>, except that the included character set name is <character set name> and, if <collate clause> is specified, then the included name of the default collation is the <collation name> contained in <collate clause>.
- 4) A privilege descriptor is created that defines the USAGE privilege on this character set to be the <authorization identifier> of the <schema definition> or <SQL-client module definition> in which the <character set definition> appears. The grantor of the privilege descriptor is set to the special grantor value “\_SYSTEM”. This privilege is grantable.

## Conformance Rules

- 1) Without Feature F451, “Character set definition”, conforming SQL language shall not contain a <character set definition>.

## 11.33 <drop character set statement>

*This Subclause is modified by Subclause 9.14, “<drop character set statement>”, in ISO/IEC 9075-4.*

### Function

Destroy a character set.

### Format

```
<drop character set statement> ::=
 DROP CHARACTER SET <character set name>
```

### Syntax Rules

- 1) Let *C* be the character set identified by the <character set name> and let *CN* be the name of *C*.
- 2) The schema identified by the explicit or implicit schema name of *CN* shall include the descriptor of *C*.
- 3) The explicit or implicit <schema name> contained in *CN* shall not be equivalent to INFORMATION\_SCHEMA.
- 4) 04 *C* shall not be referenced in any of the following:
  - a) The data type descriptor included in any column descriptor.
  - b) The data type descriptor included in any domain descriptor.
  - c) The data type descriptor generally included in any user-defined type descriptor.
  - d) The data type descriptor included in any field descriptor.
  - e) The <query expression> of any view descriptor.
  - f) The <search condition> of any constraint descriptor.
  - g) The collation descriptor of any collation.
  - h) The transliteration descriptor of any transliteration.
  - i) The SQL routine body, the <SQL parameter declaration>s, or the <returns data type> of any routine descriptor.
  - j) The <SQL parameter declaration>s or <returns data type> of any method specification descriptor.
- 5) Let the containing schema be the schema identified by the <schema name> explicitly or implicitly contained in <character set name>.

### Access Rules

- 1) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of the character set identified by *C*. The enabled authorization identifiers shall include *A*.

## General Rules

- 1) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE USAGE ON CHARACTER
SET CN FROM A
CASCADE
```

- 2) The descriptor of *C* is destroyed.

## Conformance Rules

- 1) Without Feature F451, “Character set definition”, conforming SQL language shall not contain a <drop character set statement>.

## 11.34 <collation definition>

### Function

Define a collation.

### Format

```
<collation definition> ::=
 CREATE COLLATION <collation name> FOR <character set specification>
 FROM <existing collation name> [<pad characteristic>]

<existing collation name> ::=
 <collation name>

<pad characteristic> ::=
 NO PAD
 | PAD SPACE
```

### Syntax Rules

- 1) Let *A* be the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of the <collation name>.
- 2) If a <collation definition> is contained in a <schema definition> and if the <collation name> immediately contained in the <collation definition> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the <schema definition>.
- 3) The schema identified by the explicit or implicit schema name of the <collation name> *CN* immediately contained in <collation definition> shall not include a collation descriptor whose collation name is *CN*.
- 4) The schema identified by the explicit or implicit schema name of the <collation name> *ECN* immediately contained in <existing collation name> shall include a collation descriptor whose collation name is *ECN*.
- 5) The collation identified by *ECN* shall be a collation whose descriptor includes a character repertoire name that is equivalent to that included in the descriptor of the character set identified by <character set specification>.
- 6) If <pad characteristic> is not specified, then the <pad characteristic> of the collation identified by *ECN* is implicit.
- 7) If NO PAD is specified, then the collation is said to have the NO PAD characteristic. If PAD SPACE is specified, then the collation is said to have the PAD SPACE characteristic.

### Access Rules

- 1) If a <collation definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.
- 2) The applicable privileges for *A* shall include USAGE on *ECN*.

## General Rules

- 1) A <collation definition> defines a collation.
- 2) A privilege descriptor is created that defines the USAGE privilege on this collation for *A*. The grantor of the privilege descriptor is set to the special grantor value “\_SYSTEM”.
- 3) This privilege descriptor is grantable if and only if the USAGE privilege for *A* on the collation identified by *ECN* is grantable.
- 4) A collation descriptor is created for the defined collation.
- 5) The collation descriptor *CD* created is identical to the collation descriptor for *ECN*, except that the collation name included in *CD* is *CN* and, if <pad characteristic> is specified, then the pad characteristic included in *CD* is <pad characteristic>.

## Conformance Rules

- 1) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <collation definition>.

## 11.35 <drop collation statement>

*This Subclause is modified by Subclause 9.15, “<drop collation statement>”, in ISO/IEC 9075-4.*

### Function

Destroy a collation.

### Format

```
<drop collation statement> ::=
 DROP COLLATION <collation name> <drop behavior>
```

### Syntax Rules

- 1) Let *C* be the collation identified by the <collation name> and let *CN* be the name of *C*.
- 2) The schema identified by the explicit or implicit schema name of *CN* shall include the descriptor of *C*.
- 3) The explicit or implicit <schema name> contained in *CN* shall not be equivalent to INFORMATION\_SCHEMA.
- 4) 04 If RESTRICT is specified, then *C* shall not be referenced in any of the following:
  - a) Any character set descriptor.
  - b) The triggered action of any trigger descriptor.
  - c) The <query expression> of any view descriptor.
  - d) The <search condition> of any constraint descriptor.
  - e) The SQL routine body, the <SQL parameter declaration>s, or the <returns data type> of any routine descriptor.
  - f) The <SQL parameter declaration>s or the <returns data type> of any method specification descriptor.
- 5) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of the collation identified by *C*.
- 6) Let the containing schema be the schema identified by the <schema name> explicitly or implicitly contained in <collation name>.

### Access Rules

- 1) The enabled authorization identifiers shall include *A*.



## General Rules

- 1) For every character set descriptor *CSD* that includes *CN*, *CSD* is modified such that it does not include *CN*. If *CSD* does not include any collation name, then *CSD* is modified to indicate that it utilizes the default collation for its character repertoire.
- 2) For every data type descriptor *DD* that includes *CN*, *DD* is modified such that it includes the collation name of the character set collation of the character set of *DD*.

NOTE 315 — This causes the column, domain, attribute, or field described by *DD* to revert to the default collation for its character set.

- 3) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE USAGE ON COLLATION
CN FROM A CASCADE
```

- 4) The descriptor of *C* is destroyed.

## Conformance Rules

- 1) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <drop collation statement>.

## 11.36 <transliteration definition>

### Function

Define a character transliteration.

### Format

```

<transliteration definition> ::=
 CREATE TRANSLATION <transliteration name> FOR <source character set specification>
 TO <target character set specification> FROM <transliteration source>

<source character set specification> ::=
 <character set specification>

<target character set specification> ::=
 <character set specification>

<transliteration source> ::=
 <existing transliteration name>
 | <transliteration routine>

<existing transliteration name> ::=
 <transliteration name>

<transliteration routine> ::=
 <specific routine designator>

```

### Syntax Rules

- 1) If a <transliteration definition> is contained in a <schema definition> and if the <transliteration name> immediately contained in the <transliteration definition> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the <schema definition>.
- 2) The schema identified by the explicit or implicit schema name of the <transliteration name> *TN* immediately contained in <transliteration definition> shall not include a transliteration descriptor whose transliteration name is *TN*.
- 3) The schema identified by the explicit or implicit schema name of the <character set name> *SCSN* contained in the <character set specification> contained in <source character set specification> shall include a character set descriptor whose character set name is *SCSN*.
- 4) The schema identified by the explicit or implicit schema name of the <character set name> *TCSN* contained in the <character set specification> contained in <target character set specification> shall include a character set descriptor whose character set name is *TCSN*.
- 5) If <existing transliteration name> *ETN* is specified, then:
  - a) The schema identified by the explicit or implicit schema name of *ETN* shall include a transliteration descriptor whose transliteration name is *ETN*.
  - b) The character set identified by *SCSN* shall have the same character repertoire and character encoding form as the source character set of the transliteration identified by *ETN*.

- c) The character set identified by *TCSN* shall have the same character repertoire and character encoding form as the target character set of the transliteration identified by *ETN*.
- 6) If <transliteration routine> is specified, then:
  - a) The schema identified by the explicit or implicit schema name of the <specific routine designator> *SRD* contained in <transliteration routine> shall include a routine descriptor that identifies a routine having a <specific routine designator> *SRD*.
  - b) The routine identified by *SRD* shall be an SQL-invoked function that has one parameter whose data type is character string and whose character set is the character set specified by *SCSN*; the <returns type> of the routine shall be character string whose character set is the character set specified by *TCSN*.

## Access Rules

- 1) Let *A* be the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of <transliteration name>. If a <transliteration definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.
- 2) If <existing transliteration name> is specified, then the applicable privileges for *A* shall include USAGE on the transliteration identified by *ETN*.
- 3) If <transliteration routine> is specified, then the applicable privileges for *A* shall include EXECUTE on the routine identified by *SRD*.

## General Rules

- 1) A <transliteration definition> defines a transliteration.
- 2) If <transliteration source> contains <existing transliteration name>, then let *SRDN* be the specific name included in the transliteration descriptor whose transliteration name is *TN*; otherwise, let *SRDN* be the specific name of the SQL-invoked routine identified by <transliteration routine>.
- 3) A transliteration descriptor is created that includes:
  - a) The name of the transliteration *TN*.
  - b) The name of the character set *SCSN* from which it translates.
  - c) The name of the character set *TCSN* to which it translates.
  - d) *SRDN*, the specific name of the SQL-invoked routine that performs the transliteration.
- 4) A privilege descriptor *PD* is created that defines the USAGE privilege on this transliteration to the <authorization identifier> of the <schema definition> or <SQL-client module definition> in which the <transliteration definition> appears. The grantor of the privilege descriptor is set to the special grantor value “\_SYSTEM”.
- 5) *PD* is grantable if and only if the USAGE privilege for the <authorization identifier> of the <schema definition> or <SQL-client module definition> in which the <transliteration definition> appears is also grantable on every character set identified by a <character set name> contained in the <transliteration definition>.

## **Conformance Rules**

- 1) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transliteration definition>.

## 11.37 <drop transliteration statement>

*This Subclause is modified by Subclause 9.16, “<drop transliteration statement>”, in ISO/IEC 9075-4.*

### Function

Destroy a character transliteration.

### Format

```
<drop transliteration statement> ::=
 DROP TRANSLATION <transliteration name>
```

### Syntax Rules

- 1) Let  $T$  be the transliteration identified by the <transliteration name> and let  $TN$  be the name of  $T$ .
- 2) Let  $A$  be the <authorization identifier> that owns the schema identified by the <schema name> of the transliteration identified by  $T$ .
- 3) 04 The schema identified by the explicit or implicit schema name of  $TN$  shall include the descriptor of  $T$ .
- 4)  $T$  shall not be referenced in any of the following:
  - a) The triggered action of any trigger descriptor.
  - b) The <query expression> of any view descriptor.
  - c) The <search condition> of any constraint descriptor.
  - d) The collation descriptor of any collation.
  - e) The transliteration descriptor of any translation.
  - f) The SQL routine body of any routine descriptor.

### Access Rules

- 1) The enabled authorization identifiers shall include  $A$ .

### General Rules

- 1) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE USAGE ON TRANSLATION TN FROM
A CASCADE
```

- 2) The descriptor of  $T$  is destroyed.

## **Conformance Rules**

- 1) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <drop transliteration statement>.

## 11.38 <assertion definition>

*This Subclause is modified by Subclause 9.17, “<assertion definition>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 11.8, “<assertion definition>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 12.4, “<assertion definition>”, in ISO/IEC 9075-14.*

### Function

Specify an integrity constraint.

### Format

```
14 <assertion definition> ::=
 CREATE ASSERTION <constraint name>
 CHECK <left paren> <search condition> <right paren>
 [<constraint characteristics>]
```

### Syntax Rules

- 1) If an <assertion definition> is contained in a <schema definition> and if the <constraint name> contains a <schema name>, then that <schema name> shall be equivalent to the explicit or implicit <schema name> of the containing <schema definition>.
- 2) The schema identified by the explicit or implicit schema name of the <constraint name> shall not include a constraint descriptor whose constraint name is <constraint name>.
- 3) If <constraint characteristics> is not specified, then INITIALLY IMMEDIATE NOT DEFERRABLE is implicit.
- 4) 04 The <search condition> shall not contain a <host parameter name>, an <SQL parameter name>, an <embedded variable specification> or a <dynamic parameter specification>.

NOTE 316 — <SQL parameter name> is excluded because of the scoping rules for <SQL parameter name>.

- 5) No <query expression> in the <search condition> shall reference a temporary table.
- 6) The <search condition> shall simply contain a <boolean value expression> that is retrospectively deterministic.

NOTE 317 — “retrospectively deterministic” is defined in Subclause 6.34, “<boolean value expression>”.

- 7) The <search condition> shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.

### Access Rules

- 1) Let A be the <authorization identifier> that owns the schema identified by the <schema name> of the <assertion definition>. If an <assertion definition> is contained in an <SQL-client module definition>, then the enabled authorization identifier shall include A.

## General Rules

- 1) An <assertion definition> defines an assertion. An assertion is a constraint.
  - 2) Let *SC* be the <search condition> simply contained in the <assertion definition>.
  - 3) An assertion descriptor is created that describes the assertion being defined. This descriptor includes:
    - a) The <constraint name>.
    - b) Whether the constraint is deferrable or not deferrable, as specified in <constraint characteristics>.
    - c) The initial constraint mode, as specified in <constraint characteristics>.
    - d) The applicable <search condition> *SC*.
  - 4) If the character representation of *SC* cannot be represented in the Information Schema without truncation, then a completion condition is raised: *warning — search condition too long for information schema*.
- NOTE 318 — The Information Schema is defined in [ISO9075-11].
- 5) If *SC* causes some column *CN* to be known not nullable and no other constraint causes *CN* to be known not nullable, then the nullability characteristic of *CN* is changed to known not nullable.

NOTE 319 — The nullability characteristic of a column is defined in [Subclause 4.13, “Columns, fields, and attributes”](#).

## Conformance Rules

- 1) Without Feature F521, “Assertions”, conforming SQL language shall not contain an <assertion definition>.
- 2) Without Feature F672, “Retrospective check constraints”, conforming SQL language shall not contain an <assertion definition> that generally contains CURRENT\_DATE, CURRENT\_TIMESTAMP, or LOCALTIMESTAMP.



## 11.39 <drop assertion statement>

This Subclause is modified by *Subclause 9.18*, “<drop assertion statement>”, in ISO/IEC 9075-4.

### Function

Destroy an assertion.

### Format

```
<drop assertion statement> ::=
 DROP ASSERTION <constraint name> [<drop behavior>]
```

### Syntax Rules

- 1) Let *A* be the assertion identified by <constraint name> and let *AN* be the name of *A*.
- 2) The schema identified by the explicit or implicit schema name of *AN* shall include the descriptor of *A*.
- 3) If <drop behavior> is not specified, then RESTRICT is implicit.
- 4) If RESTRICT is specified or implied, then *AN* shall not be referenced in the SQL routine body of any routine descriptor.
- 5) If *QS* is a <query specification> that contains a column reference to a column *C* in its <select list> that is not contained in a <set function specification>, and if *G* is the set of columns defined by the <grouping column reference list> of *QS*, and if the assertion *A* is needed to conclude that  $G \mapsto C$  is a known functional dependency in *QS*, then *QS* is said to be *dependent on A*.
- 6) If *V* is a view that contains a <query specification> that is dependent on *A*, then *V* is said to be *dependent on A*.
- 7) If *R* is an SQL routine whose <SQL routine body> contains a <query specification> that is dependent on *A*, then *R* is said to be *dependent on A*.
- 8) If *C* is a constraint or assertion whose <search condition> contains a <query specification> that is dependent on *A*, then *C* is said to be *dependent on A*.
- 9) If *T* is a trigger whose triggered action contains a <query specification> that is dependent on *A*, then *T* is said to be *dependent on A*.
- 10) If RESTRICT is specified or implicit, or <drop behavior> is not specified, then:
  - a) No table constraint shall be dependent on *A*.
  - b) No view shall be dependent on *A*.
  - c) No SQL routine shall be dependent on *A*.
  - d) No constraint or assertion shall be dependent on *A*.
  - e) No trigger shall be dependent on *A*.

NOTE 320 — If CASCADE is specified, then any such dependent object will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the <schema name> of the assertion identified by *AN*.

## General Rules

- 1) 04 Let *R* be any SQL-invoked routine whose routine descriptor contains the <constraint name> of *A* in the SQL routine body. Let *SN* be the <specific name> of *R*. The following <drop routine statement> is effectively executed without further Access Rule checking:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

- 2) Let *VN* be the table name of any view *V* that is dependent on *A*. The following <drop view statement> is effectively executed for every *V*:

```
DROP VIEW VN CASCADE
```

- 3) Let *SN* be the specific name of any SQL routine *SR* that is dependent on *A*, or that contains a reference to *A*. The following <drop routine statement> is effectively executed for every *SR*:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

- 4) Let *CN* be the constraint name of any constraint *C* that is dependent on *A*. Let *TN* be the name of the table constrained by *C*. The following <alter table statement> is effectively executed for every *C*:

```
ALTER TABLE TN DROP CONSTRAINT CN CASCADE
```

- 5) Let *AN2* be the constraint name of any assertion *A2* that is dependent on *A*. The following <drop assertion statement> is effectively executed for every *A2*:

```
DROP ASSERTION AN2 CASCADE
```

- 6) Let *TN* be the trigger name of any trigger *T* that is dependent on *A*. The following <drop trigger statement> is effectively executed for every *T*:

```
DROP TRIGGER TN
```

- 7) Let *SC* be the <search condition> included in the descriptor of *A*. If *SC* causes some column *CN* be to known not nullable and no other constraint causes *CN* to be known not nullable, then the nullability characteristic of *CN* is changed to possibly nullable.

NOTE 321 — The nullability characteristic of a column is defined in Subclause 4.13, “Columns, fields, and attributes”.

- 8) The descriptor of *A* is destroyed.

## Conformance Rules

- 1) Without Feature F521, “Assertions”, conforming SQL language shall not contain a <drop assertion statement>.

## 11.40 <trigger definition>

*This Subclause is modified by Subclause 9.19, “<trigger definition>”, in ISO/IEC 9075-4.*

### Function

Define triggered SQL-statements.

### Format

```

<trigger definition> ::=
 CREATE TRIGGER <trigger name> <trigger action time> <trigger event>
 ON <table name> [REFERENCING <transition table or variable list>]
 <triggered action>

<trigger action time> ::=
 BEFORE
 | AFTER
 | INSTEAD OF

<trigger event> ::=
 INSERT
 | DELETE
 | UPDATE [OF <trigger column list>]

<trigger column list> ::=
 <column name list>

<triggered action> ::=
 [FOR EACH { ROW | STATEMENT }]
 [<triggered when clause>]
 <triggered SQL statement>

<triggered when clause> ::=
 WHEN <left paren> <search condition> <right paren>

04 <triggered SQL statement> ::=
 <SQL procedure statement>
 | BEGIN ATOMIC { <SQL procedure statement> <semicolon> }... END

<transition table or variable list> ::=
 <transition table or variable>...

<transition table or variable> ::=
 OLD [ROW] [AS] <old transition variable name>
 | NEW [ROW] [AS] <new transition variable name>
 | OLD TABLE [AS] <old transition table name>
 | NEW TABLE [AS] <new transition table name>

<old transition table name> ::=
 <transition table name>

<new transition table name> ::=
 <transition table name>

<transition table name> ::=

```

```
<identifier>

<old transition variable name> ::=
 <correlation name>

<new transition variable name> ::=
 <correlation name>
```

## Syntax Rules

- 1) Case:
  - a) If a <trigger definition> is contained in a <schema definition> and if the <trigger name> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the containing <schema definition>.
  - b) If a <trigger definition> is contained in an <SQL-client module definition> and if the <trigger name> contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the <SQL-client module definition>.
- 2) Let *TN* be the <table name> of a <trigger definition>. The table *T* identified by *TN* is the *subject table* of the <trigger definition>.
- 3) The schema identified by the explicit or implicit <schema name> of *TN* shall include the descriptor of *T*.
- 4) The schema identified by the explicit or implicit <schema name> of a <trigger name> *TRN* shall not include a trigger descriptor whose trigger name is *TRN*.
- 5) Case:
  - a) If INSTEAD OF is specified, then *T* shall be a viewed table.  
Case:
    - i) If <trigger event> specifies INSERT, then *T* shall not be trigger insertable-into.
    - ii) If <trigger event> specifies UPDATE, then *T* shall not be trigger updatable.
    - iii) If <trigger event> specifies DELETE, then *T* shall not be trigger deletable.
  - b) Otherwise, *T* shall be a base table that is not a declared local temporary table.
- 6) If a <trigger column list> is specified, then:
  - a) No <column name> shall appear more than once in the <trigger column list>.
  - b) The <column name>s of the <trigger column list> shall identify columns of *T*.
- 7) If REFERENCING is specified, then:
  - a) Let *OR*, *OT*, *NR*, and *NT* be the <old transition variable name>, <old transition table name>, <new transition variable name>, and <new transition table name>, respectively.
  - b) OLD or OLD ROW, NEW or NEW ROW, OLD TABLE, and NEW TABLE shall be specified at most once each within the <transition table or variable list>.
  - c) Case:

- i) If <trigger event> specifies INSERT, then neither OLD ROW nor OLD TABLE shall be specified.
- ii) If <trigger event> specifies DELETE, then neither NEW ROW nor NEW TABLE shall be specified.
- d) No two of *OR*, *OT*, *NR*, and *NT* shall be equivalent.
- e) Both *OR* and *NR* are range variables.  
 NOTE 322 — “range variable” is defined in Subclause 4.14.6, “Operations involving tables”.
- f) The scope of *OR*, *OT*, *NR*, and *NT* is the <triggered action>, excluding any <SQL schema statement>s that are contained in the <triggered action>.
- 8) If neither FOR EACH ROW nor FOR EACH STATEMENT is specified, then FOR EACH STATEMENT is implicit.
- 9) If *OR* or *NR* is specified, then FOR EACH ROW shall be specified.
- 10) The <triggered action> shall not contain an <SQL parameter reference>, a <host parameter name>, a <dynamic parameter specification>, or an <embedded variable name>.
- 11) It is implementation-defined whether the <triggered SQL statement> shall not generally contain an <SQL transaction statement>, an <SQL connection statement>, an <SQL schema statement>, an <SQL dynamic statement>, or an <SQL session statement>.
- 12) If BEFORE is specified, then:
  - a) It is implementation-defined whether the <triggered action> shall not generally contain an <SQL data change statement> or a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.
  - b) Neither OLD TABLE nor NEW TABLE shall be specified.
  - c) The <triggered action> shall not contain a <field reference> that references a field in the new transition variable corresponding to a generated column of *T*.
- 13) If INSTEAD OF is specified, then:
  - a) <triggered when clause> shall not be specified.
  - b) <trigger column list> shall not be specified.
  - c) <table name> shall not identify a view that is either recursive or referenceable, or whose view descriptor includes an indication that WITH CASCADED CHECK OPTION has been specified.
  - d) <table name> shall not identify a view that is a leaf underlying table of the <query expression> of a view whose view descriptor includes an indication that WITH CASCADED CHECK OPTION has been specified.

## Access Rules

- 1) Let *A* be the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of the <trigger name> of the <trigger definition>. If a <trigger definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.
- 2) The applicable privileges for *A* for *T* shall include TRIGGER.

- 3) 04 If the <triggered action> *TA* of a <trigger definition> contains an <old transition table name> *OTTN*, an <old transition variable name> *OTVN*, a <new transition table name> *NTTN*, or a <new transition variable name> *NTVN*, then the applicable privileges for *T* shall include SELECT on *T*.

## General Rules

- 1) A <trigger definition> defines a trigger.
- 2) *OT* identifies the old transition table. *NT* identifies the new transition table. *OR* identifies the old transition variable. *NR* identifies the new transition variable.

NOTE 323 — “old transition table”, “new transition table”, “old transition variable”, and “new transition variable” are defined in Subclause 4.38.1, “General description of triggers”.

- 3) The transition table identified by *OT* is the table associated with *OR*. The transition table identified by *NT* is the table associated with *NR*.
- 4) If the character representation of the <triggered SQL statement> cannot be represented in the Information Schema without truncation, then a completion condition is raised: *warning — statement too long for information schema*.

NOTE 324 — The Information Schema is defined in [ISO9075-11].

- 5) A trigger descriptor is created for <trigger definition>s as follows:
- a) The trigger name included in the trigger descriptor is <trigger name>.
  - b) The subject table included in the trigger descriptor is <table name>.
  - c) The trigger action time included in the trigger descriptor is <trigger action time>.
  - d) If FOR EACH STATEMENT is specified or implicit, then an indication that the trigger is a statement-level trigger is included in the trigger descriptor; otherwise, an indication that the trigger is a row-level trigger is included in the trigger descriptor.
  - e) The trigger event included in the trigger descriptor is <trigger event>.
  - f) Any <old transition variable name>, <new transition variable name>, <old transition table name>, or <new transition table name> specified in the <trigger definition> is included in the trigger descriptor as the old transition variable name, new transition variable name, old transition table name, or new transition table name, respectively.
  - g) The triggered action included in the trigger descriptor is the specified <triggered action>.
  - h) If a <trigger column list> *TCL* is specified, then *TCL* is the trigger column list included in the trigger descriptor; otherwise, that trigger column list is empty.
  - i) The *triggered action column set* included in the trigger descriptor is the set of all distinct, fully qualified names of columns contained in the <triggered action>.
  - j) The timestamp of creation included in the trigger descriptor is the timestamp of creation of the trigger.
- 6) If INSTEAD OF is specified, then:
- a) Let *VD* be the view descriptor of the view *V* identified by the subject table of *TR*. Let *VN* be the <table name> included in *VD*. Let *QE* be the <query expression> included in *VD*. Let *U* be the <authorization identifier> that owns the schema identified by the <schema name> of *VN*.

- b) Case:
  - i) If the <trigger event> is INSERT, then
    - 1) *VD* is modified to include an indication that *V* is trigger insertable-into.
    - 2) A privilege descriptor is created that defines the privilege INSERT on *V* to *U*. This privilege is grantable. The grantor of that privilege descriptor is set to the special grantor value “\_SYSTEM”.
    - 3) For each column *C* of *V*, a privilege descriptor is created that defines the privilege INSERT on *C* to *U*. This privilege is grantable. The grantor of that privilege descriptor is set to the special grantor value “\_SYSTEM”.
  - ii) If the <trigger event> is UPDATE, then
    - 1) *VD* is modified to include an indication that *V* is trigger updatable.
    - 2) A privilege descriptor is created that defines the privilege UPDATE on *V* to *U*. This privilege is grantable. The grantor of that privilege descriptor is set to the special grantor value “\_SYSTEM”.
    - 3) For each column *C* of *V*, a privilege descriptor is created that defines the privilege UPDATE on *C* to *U*. This privilege is grantable. The grantor of that privilege descriptor is set to the special grantor value “\_SYSTEM”.
  - iii) If the <trigger event> is DELETE, then
    - 1) *VD* is modified to include an indication that *V* is trigger deletable.
    - 2) A privilege descriptor is created that defines the privilege DELETE on *V* to *U*. This privilege is grantable. The grantor of that privilege descriptor is set to the special grantor value “\_SYSTEM”.

## Conformance Rules

- 1) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain a <trigger definition>.
- 2) Without Feature T212, “Enhanced trigger capability”, in conforming SQL language, a <triggered action> shall contain FOR EACH ROW.
- 3) Without Feature T213, “INSTEAD OF triggers”, in conforming SQL language, a <trigger action time> shall not immediately contain INSTEAD OF.



## 11.41 <drop trigger statement>

### Function

Destroy a trigger.

### Format

```
<drop trigger statement> ::=
 DROP TRIGGER <trigger name>
```

### Syntax Rules

*None.*

### Access Rules

- 1) Let *TR* be the trigger identified by the <trigger name>. Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of *TR*. The enabled authorization identifiers shall include *A*.

### General Rules

- 1) If *TR* is an INSTEAD OF trigger, then:
  - a) Let *VD* be the view descriptor of the view *V* identified by the subject table of *TR*. Let *VN* be the <table name> included in *VD*. Let *QE* be the <query expression> included in *VD*. Let *U* be the <authorization identifier> that owns the schema identified by the <schema name> of *VN*.
  - b) Case:
    - i) If *TR* is an insert INSTEAD OF trigger, then:
      - 1) *VD* is modified to include an indication that *V* is not trigger insertable-into.
      - 2) Case:
        - A) If *V* is not insertable-into, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:  
REVOKE INSERT ON *VN* FROM *U* CASCADE
        - B) Otherwise:
          - I) The General Rules of [Subclause 10.12, “Determination of view component privileges”](#), are applied with *V* as *VIEW*, yielding a set of view component privilege descriptors returned from *DESCRIPTOR SET*.
          - II) Case:
            - 1) If there is no view component table privilege descriptor whose identified object is *QE*, whose action is INSERT, whose grantor is “\_SYSTEM”, and

whose grantee is *U*, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE INSERT ON VN FROM U CASCADE
```

- 2) Otherwise, let *VCTPD* be the view component table privilege descriptor whose identified object is *QE*, whose action is INSERT, whose grantor is “\_SYSTEM”, and whose grantee is *A*. If *VCTPD* indicates that the privilege is not grantable, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE GRANT OPTION FOR INSERT ON VN FROM U CASCADE
```

- ii) If *TR* is an update INSTEAD OF trigger, then:

- 1) *VD* is modified to include an indication that *V* is not trigger updatable.

- 2) Case:

- A) If *V* is not updatable, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE UPDATE ON VN FROM U CASCADE
```

- B) Otherwise:

- I) The General Rules of [Subclause 10.12, “Determination of view component privileges”](#), are applied with *V* as *VIEW*, yielding a set of view component privilege descriptors returned from *DESCRIPTOR SET*.

- II) Case:

- 1) If there is no view component table privilege descriptor whose identified object is *QE*, whose action is UPDATE, whose grantor is “\_SYSTEM”, and whose grantee is *U*, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE UPDATE ON VN FROM U CASCADE
```

- 2) Otherwise, let *VCTPD* be the view component table privilege descriptor whose identified object is *QE*, whose action is UPDATE, whose grantor is “\_SYSTEM”, and whose grantee is *A*. If *VCTPD* indicates that the privilege is not grantable, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE GRANT OPTION FOR UPDATE ON VN FROM U CASCADE
```

- iii) If *TR* is a delete INSTEAD OF trigger, then:

- 1) *VD* is modified to include an indication that *V* is not trigger deletable.

- 2) Case:

- A) If *V* is not updatable, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

REVOKE DELETE ON *VN* FROM *U* CASCADE

- B) Otherwise:

- I) The General Rules of Subclause 10.12, “Determination of view component privileges”, are applied with *V* as *VIEW*, yielding a set of view component privilege descriptors returned from *DESCRIPTOR SET*.

- II) Case:

- 1) If there is no view component table privilege descriptor whose identified object is *QE*, whose action is DELETE, whose grantor is “\_SYSTEM”, and whose grantee is *U*, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

REVOKE DELETE ON *VN* FROM *U* CASCADE

- 2) Otherwise, let *VCTPD* be the view component table privilege descriptor whose identified object is *QE*, whose action is DELETE, whose grantor is “\_SYSTEM”, and whose grantee is *A*. If *VCTPD* indicates that the privilege is not grantable, then the following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

REVOKE GRANT OPTION FOR DELETE ON *VN* FROM *U* CASCADE

- 2) The descriptor of *TR* is destroyed.

## Conformance Rules

- 1) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain a <drop trigger statement>.

## 11.42 <user-defined type definition>

*This Subclause is modified by Subclause 11.9, “<user-defined type definition>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 9.4, “<user-defined type definition>”, in ISO/IEC 9075-13.*

*This Subclause is modified by Subclause 12.5, “<user-defined type definition>”, in ISO/IEC 9075-14.*

### Function

Define a user-defined type.

### Format

```

<user-defined type definition> ::=
 CREATE TYPE <user-defined type body>

13 <user-defined type body> ::=
 <schema-resolved user-defined type name>
 [<subtype clause>]
 [AS <representation>]
 [<user-defined type option list>]
 [<method specification list>]

<user-defined type option list> ::=
 <user-defined type option> [<user-defined type option>...]

<user-defined type option> ::=
 <instantiateable clause>
 | <finality>
 | <reference type specification>
 | <cast to ref>
 | <cast to type>
 | <cast to distinct>
 | <cast to source>

<subtype clause> ::=
 UNDER <supertype name>

<supertype name> ::=
 <path-resolved user-defined type name>

<representation> ::=
 <predefined type>
 | <member list>

<member list> ::=
 <left paren> <member> [{ <comma> <member> }...] <right paren>

<member> ::=
 <attribute definition>

<instantiateable clause> ::=
 INSTANTIABLE
 | NOT INSTANTIABLE

<finality> ::=

```

```
 FINAL
 | NOT FINAL

<reference type specification> ::=
 <user-defined representation>
 | <derived representation>
 | <system-generated representation>

<user-defined representation> ::=
 REF USING <predefined type>

<derived representation> ::=
 REF FROM <list of attributes>

<system-generated representation> ::=
 REF IS SYSTEM GENERATED

<cast to ref> ::=
 CAST <left paren> SOURCE AS REF <right paren> WITH <cast to ref identifier>

<cast to ref identifier> ::=
 <identifier>

<cast to type> ::=
 CAST <left paren> REF AS SOURCE <right paren> WITH <cast to type identifier>

<cast to type identifier> ::=
 <identifier>

<list of attributes> ::=
 <left paren> <attribute name> [{ <comma> <attribute name> }...] <right paren>

<cast to distinct> ::=
 CAST <left paren> SOURCE AS DISTINCT <right paren>
 WITH <cast to distinct identifier>

<cast to distinct identifier> ::=
 <identifier>

<cast to source> ::=
 CAST <left paren> DISTINCT AS SOURCE <right paren>
 WITH <cast to source identifier>

<cast to source identifier> ::=
 <identifier>

<method specification list> ::=
 <method specification> [{ <comma> <method specification> }...]

13 <method specification> ::=
 <original method specification>
 | <overriding method specification>

<original method specification> ::=
 <partial method specification> [SELF AS RESULT] [SELF AS LOCATOR]
 [<method characteristics>]

<overriding method specification> ::=
 OVERRIDING <partial method specification>
```

## 11.42 &lt;user-defined type definition&gt;

```

<partial method specification> ::=
 [INSTANCE | STATIC | CONSTRUCTOR]
 METHOD <method name> <SQL parameter declaration list>
 <returns clause>
 [SPECIFIC <specific method name>]

<specific method name> ::=
 [<schema name> <period>]<qualified identifier>

<method characteristics> ::=
 <method characteristic>...

13 <method characteristic> ::=
 <language clause>
 | <parameter style clause>
 | <deterministic characteristic>
 | <SQL-data access indication>
 | <null-call clause>

```

## Syntax Rules

- 1) Let *UDTD* be the <user-defined type definition>, let *UDTB* be the <user-defined type body> immediately contained in *UDTD*, let *UDTN* be the <schema-resolved user-defined type name> immediately contained in *UDTB*, let *SN* be the specified or implicit <schema name> of *UDTN*, let *SS* be the SQL-schema identified by *SN*, and let *UDT* be the data type defined by *UDTD*.
- 2) If *UDTD* is contained in a <schema definition> and *UDTN* contains a <schema name>, then that <schema name> shall be equivalent to the specified or implicit <schema name> of the containing <schema definition>.
- 3) 13 *SS* shall not include a user-defined type descriptor or a domain descriptor whose name is equivalent to *UDTN*.
- 4) None of <instantiateable clause>, <finality>, <reference type specification>, <cast to ref>, <cast to type>, <cast to distinct>, or <cast to source> shall be specified more than once.
- 5) Case:
  - a) If <representation> specifies <predefined type>, then *UDTD* defines a *distinct type*.
  - b) Otherwise, *UDTD* defines a *structured type*.
- 6) If <finality> specifies FINAL, then <instantiateable clause> shall not specify NOT INSTANTIABLE.
- 7) If *UDTD* defines a distinct type, then:
  - a) Let *PSDT* be the data type identified by <predefined type>.
 

Case:

    - i) If *PSDT* is an exact numeric type, then let *SDT* be an implementation-defined exact numeric type whose precision is equal to the precision of *PSDT* and whose scale is equal to the scale of *PSDT*.
    - ii) If *PSDT* is an approximate numeric type, then let *SDT* be an implementation-defined approximate numeric type whose precision is equal to the precision of *PSDT*.
    - iii) Otherwise, let *SDT* be *PSDT*.

- b) <instantiable clause> shall not be specified.
  - c) If <finality> is not specified, then FINAL is implicit; otherwise, FINAL shall be specified.
  - d) <subtype clause> shall not be specified.
  - e) <reference type specification> shall not be specified.
  - f) Neither <cast to ref> nor <cast to type> shall be specified.
  - g) If <cast to distinct> is specified, then let *FNUDT* be <cast to distinct identifier>; otherwise, let *FNUDT* be the <qualified identifier> of *UDTN*.
  - h) If <cast to source> is specified, then let *FNSDT* be <cast to source identifier>; otherwise, the Syntax Rules of Subclause 9.7, “Type name determination”, are applied with *SDT* as *PREDEF TYPE*, yielding *FNSDT* returned from *IDENTIFIER*.
- 8) If *UDTD* specifies a structured type, then:
- a) Neither <cast to distinct> nor <cast to source> shall be specified.
  - b) If <subtype clause> is specified, then <reference type specification> shall not be specified.
  - c) If neither <subtype clause> nor <reference type specification> are specified, then <system-generated representation> is implicit.
  - d) If <instantiable clause> is not specified, then INSTANTIABLE is implicit.
  - e) <finality> shall be specified.
  - f) The *originally-defined attributes* of *UDT* are those defined by <attribute definition>s contained in <member list>. No two originally-defined attributes of *UDT* shall have equivalent <attribute name>s.
  - g) For each <attribute definition> *ATD* contained in <member list>, let *AN* be the <attribute name> contained in *ATD* and let *DT* be the <data type> contained in *ATD*. The following <original method specification>s are implicit:

```
METHOD AN ()
 RETURNS DT
 LANGUAGE SQL
 DETERMINISTIC
 CONTAINS SQL
 RETURNS NULL ON NULL INPUT
```

This is the *original method specification* of the observer function of attribute *AN*.

```
METHOD AN (ATTR DT)
 RETURNS UDTN
 SELF AS RESULT
 LANGUAGE SQL
 DETERMINISTIC
 CONTAINS SQL
 CALLED ON NULL INPUT
```

This is the *original method specification* of the mutator function of attribute *AN*.

- h) If <user-defined representation> is specified, then:
  - i) Let *BT* be <predefined type>. *BT* is the *representation type of the referencing type of UDT*.

## 11.42 &lt;user-defined type definition&gt;

- ii) *BT* shall be exact numeric or a character string type that is not a large object string type.
- iii) If <cast to ref> is specified, then let *FNREF* be <cast to ref identifier>; otherwise, let *FNREF* be the <qualified identifier> of *UDTN*.
- iv) Case:
  - 1) If <cast to type> is specified, then let *FNTYP* be <cast to type identifier>.
  - 2) Otherwise, the Syntax Rules of Subclause 9.7, “Type name determination”, are applied with *BT* as *PREDEF TYPE*, yielding *FNTYP* returned from *IDENTIFIER*.
- i) If <derived representation> is specified, then no two <attribute name>s in <list of attributes> shall be equivalent.
- j) If <subtype clause> is specified, then:
  - i) <supertype name> shall not be equivalent to *UDTN*.
  - ii) 13 The <supertype name> immediately contained in the <subtype clause> shall identify the descriptor of some structured type *SST*. *UDT* is a direct subtype of *SST*, and *SST* is a direct supertype of *UDT*.
  - iii) The descriptor of *SST* shall not include an indication that *SST* is final.
  - iv) The inherited attributes of *UDT* are the attributes described by the attribute descriptors included in the descriptor of *SST*.
  - v) If <member list> is specified, then no <attribute name> contained in <member list> shall have an attribute name that is equivalent to the attribute name of an inherited attribute.
  - vi) If the user-defined type descriptor of *SST* indicates that the referencing type of *SST* has a user-defined representation, then let *BT* be the data type described by the data type descriptor of the representation type of the referencing type of *SST* included in the user-defined type descriptor of *SST*.
    - 1) If <cast to ref> is specified, then let *FNREF* be <cast to ref identifier>; otherwise, let *FNREF* be the <qualified identifier> of *UDTN*.
    - 2) Case:
      - A) If <cast to type> is specified, then let *FNTYP* be <cast to type identifier>.
      - B) Otherwise, the Syntax Rules of Subclause 9.7, “Type name determination”, are applied with *BT* as *PREDEF TYPE*, yielding *FNTYP* returned from *IDENTIFIER*.
- k) If <cast to ref> or <cast to type> is specified, then exactly one of the following shall be true:
  - i) <user-defined representation> is specified.
  - ii) <subtype clause> is specified and the user-defined type descriptor of the direct supertype of *UDT* indicates that the referencing type of the direct supertype of *UDT* has a user-defined representation.
- 9) 13 If <method specification list> is specified, then:
  - a) 13 Let *M* be the number of <method specification>s *MS<sub>i</sub>*, 1 (one) ≤ *i* ≤ *M*, contained in <method specification list>. Let *MN<sub>i</sub>* be the <method name> of *MS<sub>i</sub>*.



- b) For  $i$  ranging from 1 (one) to  $M$ :
- i) If  $MS_i$  does not specify INSTANCE, CONSTRUCTOR, or STATIC, then INSTANCE is implicit.
  - ii) If  $MS_i$  specifies STATIC, then:
    - 1) None of SELF AS RESULT, SELF AS LOCATOR, and OVERRIDING shall be specified.
    - 2)  $MS_i$  specifies a *static method*.
  - iii) If  $MS_i$  specifies CONSTRUCTOR, then:
    - 1) SELF AS RESULT shall be specified.
    - 2) OVERRIDING shall not be specified.
    - 3)  $MN_i$  shall be equivalent to the <qualified identifier> of  $UDTN$ .
    - 4) The <returns data type> shall specify  $UDTN$ .
    - 5)  $UDTD$  shall define a structured type.
    - 6) 13  $MS_i$  specifies an *SQL-invoked constructor method*.
  - iv) Let  $RN_i$  be  $SN.MN_i$ .
  - v) If <specific method name> is not specified, then an implementation-dependent <specific method name> whose <schema name> is equivalent to  $SN$  is implicit.
  - vi) If <specific method name> contains a <schema name>, then that <schema name> shall be equivalent to  $SN$ . If <specific method name> does not contain a <schema name>, then the <schema name> of  $SN$  is implicit.
  - vii) The schema identified by the explicit or implicit <schema name> of the <specific method name> shall not include a routine descriptor whose specific name is equivalent to <specific method name> or a user-defined type descriptor that includes a method specification descriptor whose specific method name is equivalent to <specific method name>.
  - viii) Let  $PDL_i$  be the <SQL parameter declaration list> contained in  $MS_i$ .
    - 1) No two <SQL parameter name>s contained in  $PDL_i$  shall be equivalent.
    - 2) No <SQL parameter name> contained in  $PDL_i$  shall be equivalent to SELF.
  - ix) Let  $N_i$  be the number of <SQL parameter declaration>s contained in  $MS_i$ . For every <SQL parameter declaration>  $PD_{i,j}$ ,  $1 \text{ (one)} \leq j \leq N_i$ :
    - 1)  $PD_{i,j}$  shall not contain <parameter mode>. A <parameter mode> of IN is implicit.
    - 2)  $PD_{i,j}$  shall not specify RESULT.
    - 3) <parameter type>  $PT_{i,j}$  immediately contained in  $PD_{i,j}$  shall not specify ROW.
    - 4) If  $PT_{i,j}$  simply contains <locator indication>, then:
      - A)  $MS_i$  shall not specify or imply LANGUAGE SQL.

- B) 13  $PT_{i,j}$  shall specify either binary large object type, character large object type, array type, multiset type, or user-defined type.
- x) If <returns data type>  $RT$  simply contains <locator indication>, then:
- 1) LANGUAGE SQL shall not be specified or implied.
  - 2)  $RT$  shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.
  - 3) 13 <result cast> shall not be specified.
- xi) If SELF AS RESULT is specified, then the <returns data type> shall specify  $UDTN$ .
- xii) For  $k$  ranging from  $(i+1)$  to  $M$ , at least one of the following conditions shall be false:
- 1)  $MN_i$  and the <method name> of  $MS_k$  are equivalent.
  - 2) Both  $MS_i$  and  $MS_k$  either specify CONSTRUCTOR or neither specifies CONSTRUCTOR.
  - 3)  $MS_k$  has  $N_i$  <SQL parameter declaration>s.
  - 4) The data type of  $PT_{i,j}$ ,  $1 \text{ (one)} \leq j \leq N_i$ , is compatible with  $PT_{k,j}$ .
- xiii) 13 The *unaugmented SQL parameter declaration list* of  $MS_i$  is the <SQL parameter declaration list> contained in  $MS_i$ .
- xiv) If  $MS_i$  specifies <original method specification>, then:
- 1) 13 The <method characteristics> of  $MS_i$  shall contain at most one <language clause>, at most one <parameter style clause>, at most one <deterministic characteristic>, at most one <SQL-data access indication>, and at most one <null-call clause>.
  - 2) 13 If <language clause> is not specified, then LANGUAGE SQL is implicit.
  - 3) If <deterministic characteristic> is not specified, then NOT DETERMINISTIC is implicit.
  - 4) <SQL-data access indication> shall be specified.
  - 5) If <null-call clause> is not specified, then CALLED ON NULL INPUT is implicit.
  - 6) Case:
    - A) If LANGUAGE SQL is specified or implied, then:
      - I) The <returns clause> shall not specify a <result cast>.
      - II) <SQL-data access indication> shall not specify NO SQL.
      - III) <parameter style clause> shall not be specified.
      - IV) Every <SQL parameter declaration> contained in <SQL parameter declaration list> shall contain an <SQL parameter name>.
    - B) Otherwise:
      - I) 13 If <parameter style> is not specified, then PARAMETER STYLE SQL is implicit.

- II) If a <result cast> is specified, then let  $V$  be some value of the <data type> specified in the <result cast> and let  $RT$  be the <returns data type>. The following shall be valid according to the Syntax Rules of Subclause 6.12, “<cast specification>”:

CAST (  $V$  AS  $RT$  )

- III) If <result cast from type>  $RCT$  simply contains <locator indication>, then  $RCT$  shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.

- 7) Let a *conflicting method specification*  $CMS$  be a method specification that is included in the descriptor of a proper supertype of  $UDT$ , such that the following are all true:
- A) The method names of  $CMS$  and  $MN_i$  are equivalent.
  - B)  $CMS$  and  $MS_i$  have the same number of SQL parameters  $N_i$ .
  - C) Let  $PCMS_j$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $CMS$ . Let  $PMS_{i,j}$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $MS_i$ .
  - D) For  $j$  varying from 1 (one) to  $N_i$ , the declared type of  $PCMS_j$  and the declared type of  $PMS_{i,j}$  are compatible.
  - E)  $MS_i$  does not specify CONSTRUCTOR.
  - F)  $CMS$  and  $MS_i$  either both are not static methods or one of  $CMS$  and  $MS_i$  is a static method and the other is not a static method.
- 8) There shall be no conflicting method specification.
- 9) The *augmented SQL parameter declaration list*  $NPL_i$  of  $MS_i$  is defined as follows.

Case:

- A) If  $MS_i$  specifies STATIC, then let  $NPL_i$  be:

(  $PD_{i,1}$  , . . . ,  $PD_{i,N_i}$  )

- B) If  $MS_i$  specifies SELF AS RESULT and SELF AS LOCATOR, then let  $NPL_i$  be:

( SELF  $UDTN$  RESULT AS LOCATOR,  $PD_{i,1}$  , . . . ,  $PD_{i,N_i}$  )

- C) If  $MS_i$  specifies SELF AS LOCATOR, then let  $NPL_i$  be:

( SELF  $UDTN$  AS LOCATOR,  $PD_{i,1}$  , . . . ,  $PD_{i,N_i}$  )

- D) If  $MS_i$  specifies SELF AS RESULT, then let  $NPL_i$  be:

( SELF  $UDTN$  RESULT,  $PD_{i,1}$  , . . . ,  $PD_{i,N_i}$  )

- E) Otherwise, let  $NPL_i$  be:

$$( \text{ SELF } \text{ UDTN}, PD_{i,1}, \dots, PD_{i,N_i} )$$

F) Let  $AN_i$  be the number of <SQL parameter declaration>s in  $NPL_i$ .

10) If  $MS_i$  does not specify **STATIC** or **CONSTRUCTOR**, then there shall be no SQL-invoked function  $F$  that satisfies all the following conditions:

A) The routine name of  $F$  and  $RN_i$  have equivalent <qualified identifier>s.

B) If  $F$  is not a static method, then  $F$  has  $AN_i$  SQL parameters; otherwise,  $F$  has  $(AN_i-1)$  SQL parameters.

C) The data type being defined is a proper subtype of

Case:

I) If  $F$  is not a static method, then the declared type of the first SQL parameter of  $F$ .

II) Otherwise, the user-defined type whose user-defined type descriptor includes the routine descriptor of  $F$ .

D) The declared type of the  $j$ -th SQL parameter in  $NPL_i$ ,  $2 \leq j \leq AN_i$  is compatible with

Case:

I) If  $F$  is not a static method, then the declared type of  $j$ -th SQL parameter of  $F$ .

II) Otherwise, the declared type of the  $(j-1)$ -th SQL parameter of  $F$ .

11) If  $MS_i$  specifies **STATIC**, then there shall be no SQL-invoked function  $F$  that is not a static method that satisfies all the following conditions:

A) The routine name of  $F$  and  $RN_i$  have equivalent <qualified identifier>s.

B)  $F$  has  $(AN_i+1)$  SQL parameters.

C) The data type being defined is a subtype of the declared type of the first SQL parameter of  $F$ .

D) The declared type of the  $j$ -th SQL parameter in  $F$ ,  $2 \leq j \leq (AN_i+1)$ , is compatible with the declared type of the  $(j-1)$ -th SQL parameter of  $NPL_i$ .

xv) 13 If  $MS_i$  specifies <overriding method specification>, then:

1)  $MS_i$  shall not specify **STATIC** or **CONSTRUCTOR**.

2) A <returns clause> contained in  $MS_i$  shall not specify a <result cast> or <locator indication>.

3) Let the candidate original method specification  $COMS$  be an original method specification whose descriptor is included in the descriptor of a proper supertype of the user-defined type being defined, such that the following are all true:

A) The <method name> of  $COMS$  and  $MN_i$  are equivalent.

- B)  $COMS$  and  $MS_i$  have the same number of SQL parameters  $N_i$ .
  - C) Let  $PCOMS_j$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $COMS$ . Let  $POVMS_j$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $MS_i$ .
  - D) For  $i$  varying from 1 (one) to  $N_i$ , the Syntax Rules of Subclause 9.18, “Data type identity”, are applied with the declared type of  $PCOMS_j$  and the declared type of  $POVMS_j$ .
  - E) The descriptor of  $COMS$  shall not include an indication that STATIC or CONSTRUCTOR was specified.
- 4) There shall exist exactly one  $COMS$ .
  - 5)  $COMS$  shall not be the corresponding method specification of a mutator or observer function.  
NOTE 325 — “Corresponding method specification” is defined in Subclause 11.51, “<SQL-invoked routine>”.
  - 6) For  $j$  ranging from 1 (one) to  $N_i$ , all of the following shall be true:
    - A) If  $POVMS_j$  contains an <SQL parameter name>  $PNM1$ , then  $PCOMS_j$  contains an <SQL parameter name> that is equivalent to  $PNM1$ .
    - B) If  $PCOMS_j$  contains an <SQL parameter name>  $PNM2$ , then  $POVMS_j$  contains an <SQL parameter name> that is equivalent to  $PNM2$ .
    - C) If  $POVMS_j$  contains a <locator indication>, then  $PCOMS_j$  contains a <locator indication>.
    - D) If  $PCOMS_j$  contains a <locator indication>, then  $POVMS_j$  contains a <locator indication>.
  - 7) Let  $ROVMS$  be the <returns data type> of  $MS_i$ . Let  $RCOMS$  be the <returns data type> of  $COMS$ .

Case:

- A) If  $RCOMS$  is a user-defined type, then:
  - I) Let a *candidate overriding method specification*  $COVRMS$  be a method specification that is included in the descriptor of a proper supertype of  $UDT$ , such that all of the following are true:
    - 1) The <method name> of  $COVRMS$  and  $MN_i$  are equivalent.
    - 2)  $COVRMS$  and  $MS_i$  have the same number of SQL parameters  $N_i$ .
    - 3) Let  $PCOVRMS_j$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $COVRMS$ . Let  $POVMS_j$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $MS_i$ .
    - 4) For  $j$  varying from 1 (one) to  $N_i$ , the Syntax Rules of Subclause 9.18, “Data type identity”, are applied with the declared type of  $PCOVRMS_j$  and the declared type of  $POVMS_j$ .

- II) Let *NOVMS* be the number of candidate overriding method specifications. For *i* varying from 1 (one) to *NOVMS*, *ROVMS* shall be a subtype of the <returns data type> of the *i*-th candidate overriding method specification.
- B) Otherwise, the Syntax Rules of Subclause 9.18, “Data type identity”, are applied with *RCOMS* and *ROVMS*.
- 8) The augmented SQL parameter declaration list *ASPD* of *MS<sub>i</sub>* is formed from the augmented SQL parameter declaration list of *COMS* by replacing the <data type> of the first parameter (named SELF) with *UDTN*.
- 9) There shall be no SQL-invoked function *F* that satisfies all the following conditions:
  - A) The routine name of *F* and the *RN<sub>i</sub>* have equivalent <qualified identifier>s.
  - B) *F* and *ASPD* have the same number *N* of SQL parameters.
  - C) The data type being defined is a proper subtype of the declared type of the first SQL parameter of *F*.
  - D) The declared type of *POVMS<sub>i</sub>*,  $1 \text{ (one)} \leq i \leq N$ , is compatible with the declared type of SQL parameter *P<sub>i+1</sub>* of *F*.
  - E) *F* is not an SQL-invoked method.

## Access Rules

- 1) Let *A* be the <authorization identifier> that owns *SS*. If a <user-defined type definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.
- 2) The applicable privileges for *A* shall include UNDER on the <user-defined type name> specified in <subtype clause>.

## General Rules

- 1) A user-defined type descriptor *UDTDS* that describes *UDT* is created. *UDTDS* includes:
  - a) The user-defined type name *UDTN*.
  - b) If *UDT* is a distinct type or INSTANTIABLE is specified or implicit, then an indication that *UDT* is instantiable; otherwise, an indication that *UDT* is not instantiable.
  - c) An indication of whether the user-defined type is final or not final.
  - d) An indication of whether *UDT* is a distinct type or a structured type.
  - e) If *UDT* is a distinct type, then the data type descriptor of *SDT*.
  - f) If *UDT* is a structured type, then:
    - i) For each inherited attribute *IA* of *UDT*, the attribute descriptor of *IA* and an indication that *IA* is an inherited attribute.
    - ii) For each originally-defined attribute *ODA* of *UDT*, the attribute descriptor of *ODA* and an indication that *ODA* is an originally-defined attribute.

- iii) The name of the direct supertype of *UDT*.
- iv) A transform descriptor with an empty list of groups.
- v) Case:
  - 1) If <user-defined representation> is specified, then an indication that the referencing type of *UDT* has a user-defined representation, along with the data type descriptor of the representation type of the referencing type of *UDT*.
  - 2) If <derived representation> is specified, then an indication that the referencing type of *UDT* has a derived representation, along with the attributes specified by <list of attributes>.
  - 3) Otherwise, an indication that the referencing type of *UDT* has a system-defined representation.
- vi) If <subtype clause> is specified, then let *SUDT* be the direct supertype of *UDT* and let *DSUDT* be the user-defined type descriptor of *SUDT*. Let *RUDT* be the referencing type of *UDT* and let *RSUDT* be the referencing type of *SUDT*.  
Case:
  - 1) If *DSUDT* indicates that *RSUDT* has a user-defined representation, then an indication that *RUDT* has a user-defined representation and the data type descriptor of the representation type of *RSUDT* included in *DSUDT*.
  - 2) If *DSUDT* indicates that *RSUDT* has a derived representation, then an indication that *RUDT* has a derived representation and the list of attributes included in *DSUDT*.
  - 3) If *DSUDT* indicates that *RSUDT* has a system-defined representation, then an indication that *RUDT* has a system-defined representation.
- vii) The ordering form NONE.
- viii) The ordering category STATE.
- g) If <method specification list> is specified, then for every <original method specification> *ORMS* contained in <method specification list>, a method specification descriptor that includes:
  - i) An indication that the method specification is original.
  - ii) An indication of whether STATIC or CONSTRUCTOR is specified.
  - iii) The <method name> of *ORMS*.
  - iv) The <specific method name> of *ORMS*.
  - v) The <SQL parameter declaration list> contained in *ORMS* (augmented, if STATIC is not specified in *ORMS*, to include the implicit first parameter with parameter name SELF).
  - vi) The <language name> contained in the explicit or implicit <language clause>.
  - vii) 13 The explicit or implicit <parameter style> if the <language name> is SQL.
  - viii) The <returns data type>.
  - ix) The <result cast from type>, if any.
  - x) An indication of whether the method is deterministic.

## 11.42 &lt;user-defined type definition&gt;

- xi) An indication of whether the method possibly modifies SQL-data, possibly reads SQL-data, possibly contains SQL, or does not possibly contain SQL.
  - xii) An indication of whether the method should not be invoked if any argument is the null value.
- h) If <method specification list> is specified, then for every <overriding method specification> *OVMS* contained in <method specification list>, let *DCMS* be the descriptor of the corresponding original method specification. The method specification descriptor of *OVMS* includes:
- i) An indication that the method specification is overriding.
  - ii) The <method name> of *OVMS*.
  - iii) The <specific method name> of *OVMS*.
  - iv) The <SQL parameter declaration list> contained in *OVMS* (augmented to include the implicit first parameter with parameter name SELF).
  - v) The <language name> included in *DCMS*.
  - vi) The <parameter style> included in *DCMS* (if any).
  - vii) The <returns data type> of *OVMS*.
  - viii) The <result cast from type> included in *DCMS* (if any).
  - ix) The determinism indication included in *DCMS*.
  - x) The SQL-data access indication included in *DCMS*.
  - xi) The indication included in *DCMS*, whether the method should not be invoked if any argument is the null value.
- 2) If *UDTD* specifies a distinct type, then:
- a) The degree of *UDT* is 0 (zero).
  - b) The following SQL-statements are executed without further Access Rule checking:

```

CREATE FUNCTION SN.FNUDT (SDTP SDT)
 RETURNS UDTN
 LANGUAGE SQL
 DETERMINISTIC
 RETURN RV1
CREATE FUNCTION SN.FNSDT (UDTP UDTN)
 RETURNS SDT
 LANGUAGE SQL
 DETERMINISTIC
 RETURN RV2
CREATE CAST (UDTN AS SDT)
 WITH FUNCTION FNSDT (UDTN)
 AS ASSIGNMENT
CREATE CAST (SDT AS UDTN)
 WITH FUNCTION SN.FNUDT (SDT)
 AS ASSIGNMENT
CREATE TRANSFORM FOR UDTN

```



```

FNUDT (FROM SQL WITH FUNCTION FNSDT (UDTN),
 TO SQL WITH FUNCTION SN.FNUDT(SDT))

```

where: *SN* is the explicit or implicit <schema name> of *UDTN*; *RV1* is an implementation-dependent <value expression> such that for every invocation of *SN.FNUDT* with argument value *AV1*, *RV1* evaluates to the representation of *AV1* in the data type identified by *UDTN*; *RV2* is an implementation-dependent <value expression> such that for every invocation of *SN.FNSDT* with argument value *AV2*, *RV2* evaluates to the representation of *AV2* in the data type *SDT*, and *SDTP* and *UDTP* are <SQL parameter name>s arbitrarily chosen.

c) Case:

- i) 09 If *SDT* is not a large object type, then the following SQL-statement is executed without further Access Rule checking:

```

CREATE ORDERING FOR UDTN
ORDER FULL BY
MAP WITH FUNCTION FNSDT(UDTN)
FOR UDTN

```

- ii) If *SDT* is a large object type, and the SQL implementation supports Feature T042, “Extended LOB data type support”, then the following SQL-statement is executed without further Access Rule checking:

```

CREATE ORDERING FOR UDTN
ORDER EQUALS ONLY BY
MAP WITH FUNCTION FNSDT(UDTN)
FOR UDTN

```

NOTE 326 — If *SDT* is a large object type, and the SQL implementation does not support Feature T042, “Extended LOB data type support”, then no ordering for *UDTN* is created.

3) If *UDTD* specifies a structured type, then:

- a) The degree of *UDT* is the number of attributes of *UDT*, including inherited attributes. The ordinal position of an inherited attribute is its ordinal position in the direct supertype of *UDT*. The ordinal position of an attribute that is an originally-defined attribute is the ordinal position of its corresponding <attribute definition> in <member list> plus the number of inherited attributes.
- b) If INSTANTIABLE is specified, then let *V* be a value of the most specific type *UDT* such that, for every attribute *ATT* of *UDT*, invocation of the corresponding observer function on *V* yields the default value for *ATT*. The following <SQL-invoked routine> is effectively executed:

```

CREATE FUNCTION UDTN () RETURNS UDTN
RETURN V

```

This SQL-invoked function is the *constructor function* for *UDT*.

- c) If <user-defined representation> is specified or if <subtype clause> is specified and the user-defined type descriptor of the direct supertype of *UDT* indicates that the referencing type of the direct supertype of *UDT* has a user-defined representation, then the following SQL-statements are executed without further Access Rule checking:

```

CREATE FUNCTION SN.FNREF (BTP BT)
RETURNS REF(UDTN)
LANGUAGE SQL
DETERMINISTIC

```

```

 STATIC DISPATCH
 RETURN RV1
CREATE FUNCTION SN.FNTYP (UDTNP REF(UDTN))
 RETURNS BT
 LANGUAGE SQL
 DETERMINISTIC
 STATIC DISPATCH
 RETURN RV2
CREATE CAST (BT AS REF(UDTN))
 WITH FUNCTION SN.FNREF(BT)
CREATE CAST (REF(UDTN) AS BT)
 WITH FUNCTION SN.FNTYP(REF(UDTN))

```

where: *SN* is the explicit or implicit <schema name> of *UDTN*; *RV1* is an implementation-dependent <value expression> such that for every invocation of *SN.FNREF* with argument value *AV1*, *RV1* evaluates to the representation of *AV1* in the data type identified by *REF(UDTN)*; *RV2* is an implementation-dependent <value expression> such that for every invocation of *SN.FNTYP* with argument value *AV2*, *RV2* evaluates to the representation of *AV2* in the data type *BT*; and *UDTNP* is an <SQL parameter name> arbitrarily chosen.

- 4) A privilege descriptor is created that defines the USAGE privilege on *UDT* to *A*. This privilege is grantable. The grantor for this privilege descriptor is set to the special grantor value “\_SYSTEM”.
- 5) If *UDTD* specifies a structured type, then a privilege descriptor is created that defines the UNDER privilege on *UDT* to *A*. The grantor for the privilege descriptor is set to the special grantor value “\_SYSTEM”. This privilege is grantable if and only if *A* holds the UNDER privilege on the direct supertype of *UDT* WITH GRANT OPTION.

## Conformance Rules

- 1) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <member list>.
- 2) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <instantiable clause> that contains NOT INSTANTIABLE.
- 3) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <original method specification> that immediately contains SELF AS RESULT.
- 4) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <method characteristics> that contains a <parameter style> that contains GENERAL.
- 5) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <original method specification> that contains an <SQL-data access indication> that immediately contains NO SQL.
- 6) Without Feature T571, “Array-returning external SQL-invoked functions”, conforming SQL language shall not contain a <method specification> that contains a <returns clause> that satisfies either of the following conditions:
  - a) A <result cast from type> is specified that simply contains an <array type> and does not contain a <locator indication>.
  - b) A <result cast from type> is not specified and <returns data type> simply contains an <array type> and does not contain a <locator indication>.

- 7) Without Feature T572, “Multiset-returning external SQL-invoked functions”, conforming SQL language shall not contain a <method specification> that contains a <returns clause> that satisfies either of the following conditions:
  - a) A <result cast from type> is specified that simply contains a <multiset type> and does not contain a <locator indication>.
  - b) A <result cast from type> is not specified and <returns data type> simply contains a <multiset type> and does not contain a <locator indication>.
- 8) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <reference type specification>.
- 9) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <partial method specification> that contains INSTANCE or STATIC.
- 10) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <method specification list>.
- 11) Without Feature S025, “Final structured types”, in conforming SQL language, a <user-defined type definition> that defines a structured type shall contain a <finality> that is NOT FINAL.
- 12) Without Feature S028, “Permutable UDT options list”, conforming SQL language shall not contain a <user-defined type option list> in which <instantiable clause>, if specified, <finality>, <reference type specification>, if specified, <cast to ref>, if specified, <cast to type>, if specified, <cast to distinct>, if specified, and <cast to source>, if specified, do not appear in that sequence.

## 11.43 <attribute definition>

*This Subclause is modified by Subclause 9.5, “<attribute definition>”, in ISO/IEC 9075-13.*

*This Subclause is modified by Subclause 12.6, “<attribute definition>”, in ISO/IEC 9075-14.*

### Function

Define an attribute of a structured type.

### Format

```
13 <attribute definition> ::=
 <attribute name> <data type>
 [<attribute default>]
 [<collate clause>]
```

```
<attribute default> ::=
 <default clause>
```

### Syntax Rules

- 1) 13 An <attribute definition> defines a certain component of some structured type. Let *UDT* be that structured type, let *UDTN* be its name, and let *SS* be the SQL-schema whose descriptor includes the descriptor of *UDT*.
- 2) Let *AN* be the <attribute name> contained in the <attribute definition>.
- 3) The declared type *DT* of the attribute is <data type>.
- 4) <collate clause> shall not be both specified in <data type> and immediately contained in <attribute definition>. If <collate clause> is immediately contained in <attribute definition>, then it is equivalent to specifying an equivalent <collate clause> in <data type>.
- 5) *DT* shall not be based on *UDT*.

NOTE 327 — The notion of one data type being based on another data type is defined in Subclause 4.1, “Data types”.

- 6) If *DT* is a <character string type> and does not contain a <character set specification>, then the default character set for *SS* is implicit.

### Access Rules

*None.*

### General Rules

- 1) A data type descriptor is created that describes *DT*.
- 2) Let *A* be the attribute defined by <attribute definition>.
- 3) An attribute descriptor is created that describes *A*. The attribute descriptor includes:

- a)  $AN$ , the name of the attribute.
  - b) The data type descriptor of  $DT$ .
  - c) The ordinal position of the attribute in  $UDT$ .
  - d) The implicit or explicit <attribute default>.
  - e) 13 The name  $UDTN$  of the user-defined type  $UDT$ .
- 4) 13 An SQL-invoked method  $OF$  is created whose signature and result data type are as given in the descriptor of the original method specification of the observer function of  $A$ . Let  $V$  be a value in  $UDT$ . If  $V$  is the null value, then the invocation  $V.AN()$  of  $OF$  returns the result of:

CAST (NULL AS  $DT$ )

Otherwise,  $V.AN()$  returns the value of  $A$  in  $V$ .

NOTE 328 — The original method specification of the observer function of  $A$  is defined in the Syntax Rules of Subclause 11.42, “<user-defined type definition>”.

NOTE 329 — The descriptor of  $OF$  is created under the General Rules of Subclause 11.51, “<SQL-invoked routine>”.

- 5) 13 An SQL-invoked method  $MF$  is created whose signature and result data type are as given in the descriptor of the original method specification of the mutator function of  $A$ . Let  $V$  be a value in  $UDT$  and let  $AV$  be a value in  $DT$ . If  $V$  is the null value, then the invocation  $V.AN(AV)$  of  $MF$  raises an exception condition: *data exception — null value substituted for mutator subject parameter*; otherwise, the invocation  $V.AN(AV)$  returns  $V2$  such that  $V2.AN() = AV$  and for every other observer function  $ANX$  of  $UDT$ ,  $V2.ANX() = V.ANX()$ .

NOTE 330 — The original method specification of the mutator function of  $A$  is defined in the Syntax Rules of Subclause 11.42, “<user-defined type definition>”.

NOTE 331 — The descriptor of  $MF$  is created under the General Rules of Subclause 11.51, “<SQL-invoked routine>”.

## Conformance Rules

- 1) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain an <attribute definition>.
- 2) Without Feature F692, “Extended collation support”, conforming SQL language shall not contain an <attribute definition> that immediately contains a <collate clause>.
- 3) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <attribute default>.
- 4) Without Feature S026, “Self-referencing structured types”, conforming SQL language shall not contain a <data type> simply contained in an <attribute definition> that is a <reference type> whose <referenced type> is equivalent to the <schema-resolved user-defined type name> simply contained in the <user-defined type definition> that contains <attribute definition>.

## 11.44 <alter type statement>

*This Subclause is modified by Subclause 9.6, “<alter type statement>”, in ISO/IEC 9075-13.*

### Function

Change the definition of a user-defined type.

### Format

```
<alter type statement> ::=
 ALTER TYPE <schema-resolved user-defined type name> <alter type action>

<alter type action> ::=
 <add attribute definition>
 | <drop attribute definition>
 | <add original method specification>
 | <add overriding method specification>
 | <drop method specification>
```

### Syntax Rules

- 1) 13 Let  $DN$  be the <schema-resolved user-defined type name> and let  $D$  be the data type identified by  $DN$ .
- 2) The schema identified by the explicit or implicit schema name of the <schema-resolved user-defined type name> shall include the descriptor of  $D$ . Let  $S$  be that schema.
- 3) The scope of the <schema-resolved user-defined type name> is the entire <alter type statement>.
- 4) If <alter type action> contains <add attribute definition>, <drop attribute definition>, or <add overriding method specification>, then  $D$  shall be a structured type.
- 5) Let  $A$  be the <authorization identifier> that owns the schema  $S$ .

### Access Rules

- 1) If an <alter type statement> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include  $A$ .
- 2) The applicable privileges for  $A$  shall include UNDER on each proper supertype of  $D$ .

### General Rules

- 1) The user-defined type descriptor of  $D$  is modified as specified by <alter type action>.

### Conformance Rules

- 1) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <alter type statement>.

## 11.45 <add attribute definition>

### Function

Add an attribute to a user-defined type.

### Format

```
<add attribute definition> ::=
 ADD ATTRIBUTE <attribute definition>
```

### Syntax Rules

- 1) Let *D* be the user-defined type identified by the <schema-resolved user-defined type name> immediately contained in the containing <alter type statement>. Let *SPD* be any supertype of *D*. Let *SBD* be any subtype of *D*.
- 2) Let *RD* be the reference type whose referenced type is *D*. Let *SPRD* be any supertype of *RD*. Let *SBRD* be any subtype of *RD*. Let *AD* be any collection type whose element type is *D*. Let *SPAD* be any collection type whose element type is *SPD* or *SPRD*. Let *SBAD* be any collection type whose element type is *SBD* or *SBRD*.
- 3) The declared type of a column of a base table shall not be *SPRD*, *SBRD*, *SPAD*, or *SBAD*.
- 4) The declared type of a column of a base table shall not be based on *D*.  
NOTE 332 — The notion of one data type being based on another data type is defined in Subclause 4.1, “Data types”.
- 5) *SBD* shall not be the structured type of a referenceable table.
- 6) Let *M* be the mutator function resulting from the <attribute definition>, had that <attribute definition> been simply contained in the <user-defined type definition> for *D*. There shall be no SQL-invoked routine *F* that satisfies all of the following conditions:
  - a) The routine name included in the descriptor of *F* and the <schema qualified routine name> of *M* have equivalent <qualified identifier>s.
  - b) *F* has two SQL parameters.
  - c) The declared type of the first SQL parameter of *F* is a subtype or supertype of *D*.
  - d) The declared type of the second SQL parameter of *F* is compatible with the second SQL parameter of *M*.
- 7) Let *O* be the observer function resulting from the <attribute definition>, had that <attribute definition> been simply contained in the <user-defined type definition> for *D*. There shall be no SQL-invoked routine *F* that satisfies all of the following conditions:
  - a) The <schema qualified routine name> of *O* and the routine name included in the descriptor of *F* have equivalent <qualified identifier>s.
  - b) *F* has 1 (one) SQL parameter.
  - c) The declared type of the first SQL parameter of *F* is a subtype or supertype of *D*.

## Access Rules

*None.*

## General Rules

- 1) The attribute defined by the <attribute definition> is added to *D*.
- 2) In all other respects, the specification of an <attribute definition> in an <alter type statement> has the same effect as specification of the <attribute definition> simply contained in the <user-defined type definition> for *D* would have had. In particular, the degree of *D* is increased by 1 (one) and the ordinal position of that attribute is equal to the new degree of *D* as specified in the General Rules of [Subclause 11.43](#), “<attribute definition>”.
- 3) Let *A* be the attribute defined by <attribute definition>. Let *CPA* be a copy of the descriptor of *A*, modified to include an indication that the attribute is an inherited attribute.
- 4) For each proper subtype *PSBD* of *D*:
  - a) Let *DPSBD* be the descriptor of *PSBD*, let *N* be the number of attribute descriptors included in *DPSBD*, and let *DA<sub>i</sub>*, 1 (one) ≤ *i* ≤ *N*, be the attribute descriptors included in *DPSBD*.
  - b) For every *i* between 1 (one) and *N*, if *DA<sub>i</sub>* is the descriptor of an originally-defined attribute, then increase the ordinal position included in *DA<sub>i</sub>* by 1 (one).
  - c) Include *CPA* in *DPSBD*.

## Conformance Rules

*None.*



## 11.46 <drop attribute definition>

### Function

Destroy an attribute of a user-defined type.

### Format

```
<drop attribute definition> ::=
 DROP ATTRIBUTE <attribute name> RESTRICT
```

### Syntax Rules

- 1) Let *D* be the user-defined type identified by the <schema-resolved user-defined type name> immediately contained in the containing <alter type statement>.
- 2) Let *A* be the attribute identified by the <attribute name> *AN*.
- 3) *A* shall be an attribute of *D* that is not an inherited attribute, and *A* shall not be the only attribute of *D*.
- 4) Let *SPD* be any supertype of *D*. Let *SBD* be any subtype of *D*. Let *RD* be the reference type whose referenced type is *D*. Let *SPRD* be any supertype of *RD*. Let *SBRD* be any subtype of *RD*. Let *AD* be any collection type whose element type is *D*. Let *SPAD* be any collection type whose element type is *SPD* or *SPRD*. Let *SBAD* be any collection type whose element type is *SBD* or *SBRD*.
- 5) The declared type of any column of any base table shall not be *SPRD*, *SBRD*, *SPAD*, or *SBAD*.
- 6) The declared type of any column of any base table shall not be based on *D*.  
NOTE 333 — The notion of one data type being based on another data type is defined in [Subclause 4.1, “Data types”](#).
- 7) *SBD* shall not be the structured type of a referenceable table.
- 8) Let *R1* be the mutator function and let *R2* be the observer function of *A*.
  - a) *R1* and *R2* shall not be the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any of the following:
    - i) The SQL routine body of any routine descriptor.
    - ii) The <query expression> of any view descriptor.
    - iii) The <search condition> of any constraint descriptor.
    - iv) The triggered action of any trigger descriptor.
  - b) The specific names of *R1* and *R2* shall not be included in any user-defined cast descriptor.
  - c) *R1* and *R2* shall not be the ordering function in the descriptor of any user-defined type.

### Access Rules

*None.*

## General Rules

- 1) The descriptor of *A* is removed from the descriptor of every *SBD*.
- 2) The descriptor of *A* is destroyed.
- 3) The descriptors of the mutator and observer functions of *A* are destroyed.
- 4) The degree of every *SBD* is reduced by 1 (one). The ordinal position of all attributes having an ordinal position greater than the ordinal position of *A* in *SBD* is reduced by 1 (one).

## Conformance Rules

*None.*

## 11.47 <add original method specification>

### Function

Add an original method specification to a user-defined type.

### Format

```
<add original method specification> ::=
 ADD <original method specification>
```

### Syntax Rules

- 1) Let *D* be the user-defined type identified by the <schema-resolved user-defined type name> *DN* immediately contained in the containing <alter type statement>. Let *SN* be the specified or implied <schema name> of *DN*. Let *SPD* be any supertype of *D*, if any. Let *SBD* be any subtype of *D*, if any.
- 2) Let *ORMS* and *PORMS* be the <original method specification> and its immediately contained <partial method specification>, respectively.
- 3) Let *MN*, *MPDL* and *MCH* be the <method name>, the <SQL parameter declaration list> and the <method characteristics>, respectively, that are simply contained in *ORMS*. *MPDL* is called the *unaugmented SQL parameter declaration list* of *ORMS*.
- 4) If *PORMS* does not specify INSTANCE, CONSTRUCTOR, or STATIC, then INSTANCE is implicit.
- 5) If *PORMS* specifies CONSTRUCTOR, then:
  - a) SELF AS RESULT shall be specified.
  - b) *MN* shall be equivalent to the <qualified identifier> of *DN*.
  - c) The <returns data type> shall specify *DN*.
  - d) *D* shall be a structured type.
  - e) *PORMS* specifies an *SQL-invoked constructor method*.
- 6) If *PORMS* specifies STATIC, then:
  - a) Neither SELF AS RESULT nor SELF AS LOCATOR shall be specified.
  - b) *PORMS* specifies a static method.
- 7) Let *RN* be *SN.MN*.
- 8) Case:
  - a) If *PORMS* does not specify <specific method name>, then an implementation-dependent <specific method name> is implicit whose <schema name> is equivalent to *SN*.
  - b) Otherwise,  
Case:

## 11.47 &lt;add original method specification&gt;

- i) If <specific method name> contains a <schema name>, then that <schema name> shall be equivalent to *SN*.
- ii) Otherwise, the <schema name> *SN* is implicit.

The schema identified by the explicit or implicit <schema name> of the <specific method name> shall not include a routine descriptor whose specific name is equivalent to <specific method name> or a user-defined type descriptor that includes a method specification descriptor whose specific method name is equivalent to <specific method name>.

- 9) *MCH* shall contain at most one <language clause>, at most one <parameter style clause>, at most one <deterministic characteristic>, at most one <SQL-data access indication>, and at most one <null-call clause>.
  - a) If <language clause> is not specified in *MCH*, then LANGUAGE SQL is implicit.
  - b) Case:
    - i) If LANGUAGE SQL is specified or implied, then:
      - 1) <parameter style clause> shall not be specified.
      - 2) <SQL-data access indication> shall not specify NO SQL.
      - 3) Every <SQL parameter declaration> contained in <SQL parameter declaration list> shall contain an <SQL parameter name>.
      - 4) The <returns clause> shall not specify a <result cast>.
    - ii) Otherwise:
      - 1) If <parameter style clause> is not specified, then PARAMETER STYLE SQL is implicit.
      - 2) If a <result cast> is specified, then let *V* be some value of the <data type> specified in the <result cast> and let *RT* be the <returns data type>. The following shall be valid according to the Syntax Rules of Subclause 6.12, “<cast specification>”:
 

```
CAST (V AS RT)
```
      - 3) If <result cast from type> *RCT* simply contains <locator indication>, then *RCT* shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.
  - c) If <deterministic characteristic> is not specified in *MCH*, then NOT DETERMINISTIC is implicit.
  - d) If <SQL-data access indication> is not specified, then CONTAINS SQL is implicit.
  - e) If <null-call clause> is not specified in *MCH*, then CALLED ON NULL INPUT is implicit.
- 10) No two <SQL parameter name>s contained in *MPDL* shall be equivalent.
- 11) No <SQL parameter name> contained in *MPDL* shall be equivalent to SELF.
- 12) Let *N* be the number of <SQL parameter declaration>s contained in *MPDL*. For every <SQL parameter declaration> *PD<sub>j</sub>*, 1 (one) ≤ *j* ≤ *N*:
  - a) *PD<sub>j</sub>* shall not contain <parameter mode>. A <parameter mode> of IN is implicit.

- b)  $PD_j$  shall not specify RESULT.
  - c) <parameter type>  $PT_j$  immediately contained in  $PD_j$  shall not specify ROW.
  - d) If  $PT_j$  simply contains <locator indication>, then:
    - i)  $MCH$  shall not specify LANGUAGE SQL, nor shall LANGUAGE SQL be implied.
    - ii)  $PT_j$  shall specify either binary large object type, character large object type, array type, multiset type, or user-defined type.
- 13) If <returns data type>  $RT$  simply contains <locator indication>, then:
- a)  $MCH$  shall not be specify LANGUAGE SQL, nor shall LANGUAGE SQL be implied.
  - b)  $RT$  shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.
  - c) <result cast> shall not be specified.
- 14) If SELF AS RESULT is specified, then the <returns data type> shall specify  $DN$ .
- 15) Case:
- a) If  $ORMS$  specifies CONSTRUCTOR, then let a *conflicting method specification*  $CMS$  be a method specification whose descriptor is included in the descriptor of  $D$ , such that the following are all true:
    - i)  $MPDL$  and the unaugmented SQL parameter list of  $CMS$  have the same number  $N$  of SQL parameters.
    - ii) Let  $PCMS_j$ ,  $1 \text{ (one)} \leq j \leq N$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $CMS$ . Let  $PMS_j$ ,  $1 \text{ (one)} \leq j \leq N$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list  $MPDL$ .
    - iii) For  $j$  varying from 1 (one) to  $N$ , the declared type of  $PCMS_j$  and the declared type of  $PMS_j$  are compatible.
    - iv)  $CMS$  is an SQL-invoked constructor method.
  - b) Otherwise, let a *conflicting method specification*  $CMS$  be a method specification whose descriptor is included in the descriptor of some  $SPD$  or  $SBD$ , such that the following are all true:
    - i)  $MN$  and the method name included in the descriptor of  $CMS$  are equivalent.
    - ii)  $MPDL$  and the unaugmented SQL parameter list of  $CMS$  have the same number  $N$  of SQL parameters.
    - iii) Let  $PCMS_j$ ,  $1 \text{ (one)} \leq j \leq N$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list of  $CMS$ . Let  $PMS_j$ ,  $1 \text{ (one)} \leq j \leq N$ , be the  $j$ -th SQL parameter in the unaugmented SQL parameter declaration list  $MPDL$ .
    - iv) For  $j$  varying from 1 (one) to  $N$ , the declared type of  $PCMS_j$  and the declared type of  $PMS_j$  are compatible.
    - v)  $CMS$  and  $ORMS$  either both are not instance methods or one of  $CMS$  and  $ORMS$  is a static method and the other is an instance method.

- 16) There shall be no conflicting method specification.
- 17) Let  $MP_i$ ,  $1 \text{ (one)} \leq i \leq N$ , be the  $i$ -th <SQL parameter declaration> contained in  $MPDL$ . The augmented SQL parameter declaration list  $NPL$  of  $ORMS$  is defined as follows.

Case:

- a) If  $PORMS$  specifies **STATIC**, then let  $NPL$  be:

(  $MP_1$ , ...,  $MP_N$  )

- b) If  $ORMS$  specifies **SELF AS RESULT** and **SELF AS LOCATOR**, then let  $NPL$  be:

( SELF DN RESULT AS LOCATOR,  
 $MP_1$ , ...,  
 $MP_N$  )

- c) If  $ORMS$  specifies **SELF AS LOCATOR**, then let  $NPL$  be:

( SELF DN AS LOCATOR,  
 $MP_1$ , ...,  
 $MP_N$  )

- d) If  $ORMS$  specifies **SELF AS RESULT**, then let  $NPL$  be:

( SELF DN RESULT,  $MP_1$ , ...,  
 $MP_N$  )

- e) Otherwise, let  $NPL$  be:

( SELF DN,  $MP_1$ , ...,  
 $MP_N$  )

Let  $AN$  be the number of <SQL parameter declaration>s in  $NPL$ .

- 18) If  $PORMS$  does not specify **STATIC** or **CONSTRUCTOR**, then there shall be no SQL-invoked function  $F$  that satisfies all the following conditions:
- a)  $F$  is not an SQL-invoked method.
  - b) The <routine name> of  $F$  and  $RN$  have equivalent <qualified identifier>s.
  - c)  $F$  has  $AN$  SQL parameters.
  - d)  $D$  is a subtype or supertype of the declared type of the first SQL parameter of  $F$ .
  - e) The declared type of the  $i$ -th SQL parameter in  $NPL$ ,  $2 \leq i \leq AN$  is compatible with the declared type of  $i$ -th SQL parameter of  $F$ .
- 19) If  $PORMS$  specifies **STATIC**, then there shall be no SQL-invoked function  $F$  that is not a static method that satisfies all the following conditions:
- a) The <routine name> of  $F$  and  $RN$  have equivalent <qualified identifier>s.
  - b)  $F$  has  $(AN+1)$  SQL parameters.

- c)  $D$  is a subtype or supertype of the declared type of the first SQL parameter of  $F$ .
- d) The declared type of the  $i$ -th SQL parameter of  $F$ ,  $2 \leq i \leq (AN+1)$ , is compatible with the declared type of the  $(i-1)$ -th SQL parameter of  $NPL$ .

## Access Rules

*None.*

## General Rules

- 1) Let  $STDS$  be the descriptor of  $D$ . A method specification descriptor  $DOMS$  is created for  $ORMS$ .  $DOMS$  includes:
  - a) An indication that the method specification is original.
  - b) An indication of whether **STATIC** or **CONSTRUCTOR** is specified.
  - c) The <method name>  $MN$ .
  - d) The <specific method name> contained in  $PORMS$ .
  - e) The augmented SQL parameter declaration list  $NPL$ .
  - f) For every parameter descriptor of a parameter of  $NPL$ , a locator indication (if specified).
  - g) The <returns data type> contained in  $PORMS$ .
  - h) The <result cast from type> contained in  $PORMS$  (if any).
  - i) The locator indication, if a <locator indication> is contained in the <returns clause> of  $PORMS$  (if any).
  - j) The <language name> explicitly or implicitly contained in  $MCH$ .
  - k) The explicit or implicit <parameter style> contained in  $MCH$ , if the <language name> is not **SQL**.
  - l) The determinism indication contained in  $MCH$ .
  - m) An indication of whether the method possibly modifies SQL-data, possibly reads SQL-data, possibly contains SQL, or does not possibly contain SQL.
  - n) An indication of whether the method should not be invoked if any argument is the null value.
  - o) The **CURRENT\_TIMESTAMP** as the value of the creation timestamp.
- 2)  $DOMS$  is added to  $STDS$ .
- 3) Let  $N$  be the number of table descriptors that include the user-defined type name of a subtype of  $D$ .  
For  $i$  varying from 1 (one) to  $N$ :
  - a) Let  $TN_i$  be the <table name> included in the  $i$ -th such table descriptor.

- b) For every table privilege descriptor that specifies  $TN_i$  and a privilege of SELECT, a new table/method privilege descriptor is created that specifies  $TN_i$ , the same action, grantor, and grantee, and the same grantability, and the <specific method name> contained in *ORMS*.

## **Conformance Rules**

*None.*



## 11.48 <add overriding method specification>

### Function

Add an overriding method specification to a user-defined type.

### Format

```
<add overriding method specification> ::=
 ADD <overriding method specification>
```

### Syntax Rules

- 1) Let *OVMS* be the <overriding method specification> immediately contained in <add overriding method specification>. Let *D* be the user-defined type identified by the <schema-resolved user-defined type name> *DN* immediately contained in the <alter type statement> containing *OVMS*. Let *SN* be the specified or implied <schema name> of *DN*. Let *SPD* be any supertype of *D*, if any. Let *SBD* be any subtype of *D*, if any.
- 2) Let *POVMS* be the <partial method specification> immediately contained in *OVMS*. *POVMS* shall not specify **STATIC** or **CONSTRUCTOR**.
- 3) Let *MN*, *RTC* and *MPDL* be <routine name>, the <returns clause> and the <SQL parameter declaration list> immediately contained in *POVMS*.
- 4) *MN* shall not be equivalent to the <qualified identifier> of the user-defined type name of any *SPD* or *SBD* other than *D*.
- 5) Let *RN* be *SN.MN*.
- 6) Case:
  - a) If *POVMS* does not specify <specific method name>, then an implementation-dependent <specific method name> is implicit whose <schema name> is equivalent to *SN*.
  - b) Otherwise,  
Case:
    - i) If <specific method name> contains a <schema name>, then that <schema name> shall be equivalent to *SN*.
    - ii) Otherwise, the <schema name> *SN* is implicit.
- 7) The schema identified by the explicit or implicit <schema name> of the <specific method name> shall not include a routine descriptor whose specific name is equivalent to <specific method name> or a user-defined type descriptor that includes a method specification descriptor whose specific method name is equivalent to <specific method name>.
- 8) *RTC* shall not specify a <result cast> or <locator indication>.
- 9) Let the *candidate original method specification COMS* be an original method specification that is included in the descriptor of a proper supertype of the user-defined type of *D*, such that the following are all true:

## 11.48 &lt;add overriding method specification&gt;

- a) *MN* and the <method name> of *COMS* are equivalent.
- b) Let *N* be the number of elements of the augmented SQL parameter declaration list *UPCOMS* generally included in the descriptor of *COMS*. *MPDL* contains (*N*-1) SQL parameter declarations.
- c) For *i* varying from 2 to *N*, the Syntax Rules of Subclause 9.18, “Data type identity”, are applied with the data types of the SQL parameters *PCOMS<sub>i</sub>* of *UPCOMS* and the data types of the SQL parameters *POVMS<sub>i-1</sub>* of *MPDL*, respectively.
- d) The descriptor of *COMS* shall not include an indication that STATIC or CONSTRUCTOR was specified.

10) There shall exist exactly one such *COMS*.

11) *COMS* shall not be the corresponding method specification of a mutator or observer function.

NOTE 334 — “Corresponding method specification” is defined in Subclause 11.51, “<SQL-invoked routine>”.

12) For *i* varying from 2 to *N*:

- a) If *POVMS<sub>i-1</sub>* contains an <SQL parameter name> *PNM1*, then the descriptor of the *i*-th parameter of the augmented <SQL parameter declaration list> of *UPCOMS* shall include a parameter name that is equivalent to *PNM1*.
- b) If the descriptor of the *i*-th parameter of the augmented <SQL parameter declaration list> of *UPCOMS* includes a parameter name *PNM2*, then *POVMS<sub>i-1</sub>* shall contain an <SQL parameter name> that is equivalent to *PNM2*.
- c) *POVMS<sub>i-1</sub>* shall not contain <parameter mode>. A <parameter mode> IN is implicit.
- d) *POVMS<sub>i-1</sub>* shall not specify RESULT.
- e) If the <parameter type> *PT<sub>i-1</sub>* immediately contained in *POVMS<sub>i-1</sub>* contains a <locator indication>, then the descriptor of the *i*-th parameter of the augmented <SQL parameter declaration list> of *UPCOMS* shall include a <locator indication>.
- f) If the the descriptor of the *i*-th parameter of the augmented <SQL parameter declaration list> of *UPCOMS* includes a <locator indication>, then the <parameter type> *PT<sub>i-1</sub>* immediately contained in *POVMS<sub>i-1</sub>* shall contain a <locator indication>.

13) Let *ROVMS* be the <returns data type> of *RTC*. Let *RCOMS* be the <returns data type> of *COMS*.

Case:

- a) If *RCOMS* is a user-defined type, then:
  - i) Let a *candidate overriding method specification COVRMS* be a method specification that is included in the descriptor of a proper supertype or a proper subtype of *UDT*, such that all of the following are true:
    - 1) The <method name> of *COVRMS* and *MN* are equivalent.
    - 2) *COVRMS* and *OVMS* have the same number of SQL parameters *N*.
    - 3) Let *PCOVRMS<sub>i</sub>*, 1 (one) ≤ *i* ≤ *N*, be the *i*-th SQL parameter in the unaugmented SQL parameter declaration list of *COVRMS*. Let *POVMS<sub>i</sub>*, 1 (one) ≤ *i* ≤ *N*, be the *i*-th SQL parameter in the unaugmented SQL parameter declaration list of *OVMS*.

- 4) For  $i$  varying from 1 (one) to  $N$ , the Syntax Rules of [Subclause 9.18, “Data type identity”](#), are applied with the declared type of  $PCOVRMS_i$  and the declared type of  $POVMS_i$ .
- ii) Let  $NOVMS$  be the number of candidate overriding method specifications. For  $i$  varying from 1 (one) to  $NOVMS$ , let  $COVRMS_i$  be the  $i$ -th candidate overriding method specification.

Case:

- 1) If  $COVRMS_i$  is included in the descriptor of a proper supertype of  $D$ , then  $ROVMS$  shall be a subtype of the <returns data type> of  $COVRMS_i$ .
  - 2) Otherwise,  $ROVMS$  shall be a supertype of the <returns data type> of  $COVRMS_i$ .
  - b) Otherwise, the Syntax Rules of [Subclause 9.18, “Data type identity”](#), are applied with  $RCOMS$  and  $ROVMS$  as the data types.
- 14) Let a *conflicting overriding method specification*  $COVMS$  be an overriding method specification that is included in the descriptor of  $D$ , such that all of the following are true:
- a)  $MN$  and the method name of  $COVMS$  are equivalent.
  - b) The augmented SQL parameter declaration list of  $COVMS$  contains  $N$  elements.
  - c) For  $i$  varying from 2 to  $N$ , the data types of the SQL parameter  $POVMS_{i-1}$  and the SQL parameter  $PCOVMS_{i-1}$  of  $COVMS$  are compatible.
- 15) There shall be no conflicting overriding method specification.
- 16) The augmented SQL parameter declaration list  $ASPD_L$  of  $OVMS$  is formed from the augmented SQL parameter declaration list of  $COMS$  by replacing the <data type> of the first parameter (named SELF) with the <schema-resolved user-defined type name>  $DN$ .
- 17) There shall be no SQL-invoked function  $F$  that satisfies all the following conditions:
- a)  $F$  is not an SQL-invoked method.
  - b) The <routine name> of  $F$  and the <routine name>  $MS$  have equivalent <qualified identifier>s.
  - c) Let  $NPF$  be the number of SQL parameters in  $ASPD_L$ .  $F$  has  $NPF$  SQL parameters.
  - d)  $D$  is a subtype or supertype of the declared type of the first SQL parameter of  $F$ .
  - e) The declared type of the  $i$ -th SQL parameter in  $ASPD_L$ ,  $2 \leq i \leq NPF$  is compatible with the declared type of  $i$ -th SQL parameter of  $F$ .
- 18) If the descriptor of  $D$  includes any method specification descriptor, then:
- a) Let  $M$  be the number of method specification descriptors  $MSD_i$ ,  $1 \text{ (one)} \leq i \leq M$ , included in the descriptor of  $D$ .
  - b) For  $i$  ranging from 1 (one) to  $M$ :
    - i) Let  $N_i$  be the number of <SQL parameter declaration>s contained in the augmented SQL parameter declaration list included in  $MSD_i$ . Let  $PT_{i,j}$ ,  $1 \text{ (one)} \leq j \leq N_i$ , be the  $j$ -th <parameter type> contained in  $MSD_i$ .
    - ii) At least one of the following conditions shall be false:

- 1) The <routine name> included in  $MSD_i$  is equivalent to  $MN$ .
- 2)  $ASPD_L$  has  $N_i$  <SQL parameter declaration>s.
- 3) The data type of  $PT_{i,j}$ ,  $1 \text{ (one)} \leq j \leq N_i$ , is compatible with the data type of the  $j$ -th <SQL parameter declaration> of  $MPDL$ .
- 4)  $MSD_i$  does not include an indication that CONSTRUCTOR was specified.

## Access Rules

*None.*

## General Rules

- 1) Let  $STDS$  be the descriptor of  $D$ , and let  $DCMS$  be the descriptor of the corresponding original method specification  $COMS$ . A method specification descriptor  $DOMS$  is created for  $OVMS$ .  $DOMS$  includes:
  - a) An indication that the method specification is overriding.
  - b) The <method name>  $MN$ .
  - c) The <specific method name> contained in  $POVMS$ .
  - d) The augmented SQL parameter declaration list  $APDL$ .
  - e) For every parameter descriptor of a parameter of  $APDL$ , the locator indication of the descriptor of the corresponding parameter included in  $DCMS$  (if any).
  - f) The <language name> included in  $DCMS$ .
  - g) The <parameter style> included in  $DCMS$  (if any).
  - h) The <returns data type> contained in  $POVMS$ .
  - i) The <result cast from type> included in  $DCMS$  (if any).
  - j) The locator indication contained in the <returns clause> included in the  $DCMS$ .
  - k) The determinism indication included in  $DCMS$ .
  - l) The SQL-data access indication included in  $DCMS$  (if any).
  - m) The indication included in  $DCMS$  (if at all), whether the method should be invoked if any argument is the null value.
  - n) The CURRENT\_TIMESTAMP as the value of the creation timestamp.
- 2)  $DOMS$  is added to  $STDS$ .
- 3) Let  $N$  be the number of table descriptors that include the user-defined type name of a subtype of  $D$ .  
For  $i$  varying from 1 (one) to  $N$ :
  - a) Let  $TN_i$  be the <table name> included in the  $i$ -th such table descriptor.

- b) Let  $M$  be the number of table/method privilege descriptors that specify  $TN_i$  and the <specific method name> contained in  $COMS$ . For  $j$  varying from 1 (one) to  $M$ :
  - i) Let  $TMPD_j$  be the  $j$ -th such table/method privilege descriptor.
  - ii) A new table/method privilege descriptor is created that specifies  $TN_i$ , the same action, grantor, and grantee, and the same grantability, and the <specific method name> contained in  $OVMS$ .
  - iii)  $TMPD_j$  is deleted.

## Conformance Rules

*None.*

## 11.49 <drop method specification>

### Function

Remove a method specification from a user-defined type.

### Format

```
<drop method specification> ::=
 DROP <specific method specification designator> RESTRICT

<specific method specification designator> ::=
 [INSTANCE | STATIC | CONSTRUCTOR]
 METHOD <method name> <data type list>
```

### Syntax Rules

- 1) Let *D* be the user-defined type identified by the <schema-resolved user-defined type name> *DN* immediately contained in the <alter type statement> containing the <drop method specification> *DORMS*. Let *DSN* be the explicit or implicit <schema name> of *DN*. Let *SMSD* be the <specific method specification designator> immediately contained in *DORMS*.
- 2) If *SMSD* immediately contains a <specific method name> *SMN*, then:
  - a) If *SMN* contains a <schema name>, then that <schema name> shall be equivalent to *DSN*. Otherwise, the <schema name> *DSN* is implicit.
  - b) The descriptor of *D* shall include a method specification descriptor *DOOMS* whose specific method name is equivalent to *SMN*.
  - c) Let *PDL* be the augmented parameter list included in *DOOMS*.
  - d) Let *MN* be the <method name> included in *DOOMS*.
- 3) If *SMSD* immediately contains a <method name> *ME*, then:
  - a) If none of *INSTANCE*, *STATIC*, or *CONSTRUCTOR* is immediately contained in *SMSD*, then *INSTANCE* is implicit.
  - b) The descriptor of *D* shall include a method specification descriptor *DOOMS* whose method name *MN* is equivalent to *ME*.
  - c) If *SMSD* immediately contains a <data type list> *DTL*, then
 

Case:

    - i) If *STATIC* is specified, then the descriptor of *D* shall include exactly one method specification descriptor *DOOMS* that includes:
      - 1) An indication that the method specification is *STATIC*.
      - 2) An indication that the method specification is original.

- 3) An augmented parameter list *PDL* such that the declared type of its *i*-th parameter, for all *i*, is identical to the *i*-th declared type in *DTL*.
- ii) If CONSTRUCTOR is specified, then the descriptor of *D* shall include exactly one method specification descriptor *DOOMS* that includes:
  - 1) An indication that the method specification is CONSTRUCTOR.
  - 2) An indication that the method specification is original.
  - 3) An augmented parameter list *PDL* such that the declared type of its *i*-th parameter, for all *i* > 1 (one), is identical to the (*i*–1)-th declared type in *DTL* and the declared type of the first parameter of *PDL* is identical to *DN*.
- iii) Otherwise, the descriptor of *D* shall include exactly one method specification descriptor *DOOMS* for which:
  - 1) If *DOOMS* includes an indication that the method specification is original, then *DOOMS* shall not include an indication that the method specification is either STATIC or CONSTRUCTOR.
  - 2) *DOOMS* includes an augmented parameter list *PDL* such that the declared type of its *i*-th parameter, for all *i* > 1 (one), is identical to the (*i*–1)-th declared type in *DTL* and the declared type of the first parameter of *PDL* is identical to *DN*.
- d) If *SMSD* does not immediately contain a <data type list>, then

Case:

- i) If STATIC is specified, then the descriptor of *D* shall include exactly one method specification descriptor *DOOMS* that includes an indication that the method specification is both original and STATIC.
- ii) If CONSTRUCTOR is specified, then the descriptor of *D* shall include exactly one method specification descriptor *DOOMS* that includes an indications that the method specification is both original and CONSTRUCTOR.
- iii) Otherwise, the descriptor of *D* shall include exactly one method specification descriptor *DOOMS* for which if *DOOMS* includes an indication that the method specification is original, then *DOOMS* shall not include an indication that the method specification is either STATIC or CONSTRUCTOR.

4) Case:

- a) If *DOOMS* includes an indication that the method specification is original, then

Case:

- i) If *DOOMS* includes an indication that the method specification specified STATIC, then there shall be no SQL-invoked function *F* that satisfies all of the following conditions:
  - 1) The <routine name> of *F* and *MN* have equivalent <qualified identifier>s.
  - 2) If *N* is the number of elements in *PDL*, then *F* has *N* SQL parameters.
  - 3) The declared type of the first SQL parameter of *F* is *D*.

- 4) The declared type of the  $i$ -th element of  $PDL$ ,  $1 \text{ (one)} \leq i \leq N$ , is compatible with the declared type of SQL parameter  $P_i$  of  $F$ .
  - 5)  $F$  is an SQL-invoked method.
  - 6)  $F$  includes an indication that STATIC is specified.
- ii) If  $DOOMS$  includes an indication that the method specification specified CONSTRUCTOR, then there shall be no SQL-invoked function  $F$  that satisfies all of the following conditions:
- 1) The <routine name> of  $F$  and  $MN$  have equivalent <qualified identifier>s.
  - 2) If  $N$  is the number of elements in  $PDL$ , then  $F$  has  $N$  SQL parameters.
  - 3) The declared type of the first SQL parameter of  $F$  is  $D$ .
  - 4) The declared type of the  $i$ -th element of  $PDL$ ,  $2 \leq i \leq N$ , is compatible with the declared type of SQL parameter  $P_i$  of  $F$ .
  - 5)  $F$  is an SQL-invoked method.
  - 6)  $F$  includes an indication that CONSTRUCTOR is specified.
- iii) Otherwise:
- 1) There shall be no proper subtype  $PSBD$  of  $D$  whose descriptor includes the descriptor  $DOVMS$  of an overriding method specification such that all of the following is true:
    - A)  $MN$  and the <method name> included in  $DOVMS$  have equivalent <qualified identifier>s.
    - B) If  $N$  is the number of elements in  $PDL$ , then the augmented SQL parameter declaration list  $APDL$  included in  $DOVMS$  has  $N$  SQL parameters.
    - C)  $PSBD$  is the declared type of the first SQL parameter of  $APDL$ .
    - D) The declared type of the  $i$ -th element of  $PDL$ ,  $2 \leq i \leq N$ , is compatible with the declared type of SQL parameter  $P_i$  of  $APDL$ .
  - 2) There shall be no SQL-invoked function  $F$  that satisfies all of the following conditions:
    - A) The <routine name> of  $F$  and  $MN$  have equivalent <qualified identifier>s.
    - B) If  $N$  is the number of elements in  $PDL$ , then  $F$  has  $N$  SQL parameters.
    - C) The declared type of the first SQL parameter of  $F$  is  $D$ .
    - D) The declared type of the  $i$ -th element of  $PDL$ ,  $2 \leq i \leq N$ , is compatible with the declared type of SQL parameter  $P_i$  of  $F$ .
    - E)  $F$  is an SQL-invoked method.
    - F)  $F$  does not include an indication that either STATIC or CONSTRUCTOR is specified.
- b) Otherwise, there shall be no SQL-invoked function  $F$  that satisfies all of the following conditions:
- i) The <routine name> of  $F$  and  $MN$  have equivalent <qualified identifier>s.
  - ii) If  $N$  is the number of elements in  $PDL$ , then  $F$  has  $N$  SQL parameters.



- iii) The declared type of the first SQL parameter of  $F$  is  $D$ .
- iv) The declared type of the  $i$ -th element of  $PDL$ ,  $2 \leq i \leq N$ , is compatible with the declared type of SQL parameter  $P_i$  of  $F$ .
- v)  $F$  is an SQL-invoked method.
- vi)  $F$  does not include an indication that either **STATIC** or **CONSTRUCTOR** is specified.

### Access Rules

*None.*

### General Rules

- 1) Let  $STDS$  be the descriptor of  $D$ .
- 2)  $DOOMS$  is removed from  $STDS$ .
- 3)  $DOOMS$  is destroyed.

### Conformance Rules

*None.*

## 11.50 <drop data type statement>

*This Subclause is modified by Subclause 9.23, “<drop data type statement>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 9.7, “<drop data type statement>”, in ISO/IEC 9075-13.*

### Function

Destroy a user-defined type.

### Format

```
<drop data type statement> ::=
 DROP TYPE <schema-resolved user-defined type name> <drop behavior>
```

### Syntax Rules

- 1) Let *DN* be the <schema-resolved user-defined type name> and let *D* be the data type identified by *DN*. Let *SD* be any supertype of *D*.
- 2) Let *RD* be the reference type whose referenced type is *D*. Let *SRD* be any supertype of *RD*. Let *AD* be any collection type whose element type is *D*. Let *SAD* be any collection type whose element type is a supertype of *D* or *RD*.
- 3) The schema identified by the explicit or implicit schema name of *DN* shall include the descriptor of *D*.
- 4) If RESTRICT is specified, then:
  - a) The declared type of no column, field, or attribute whose descriptor is not included in the descriptor of *D* shall be *SRD* or *SAD*.
  - b) The declared type of no column, attribute, or field shall be based on *D*.
  - c) *D* shall have no proper subtypes.
  - d) *D* shall not be the structured type of a referenceable table.
  - e) The transform descriptor included in the user-defined type descriptor of *D* shall include an empty list of transform groups.
  - f) *D*, *RD*, and *AD* shall not be referenced in any of the following:
    - i) The <query expression> of any view descriptor.
    - ii) The <search condition> of any constraint descriptor.
    - iii) A triggered action of any trigger descriptor.
    - iv) A user-defined cast descriptor.
    - v) 04 A user-defined type descriptor other than that of *D* itself.
  - g) There shall be no SQL-invoked routine that is not dependent on *D* and whose routine descriptor includes the descriptor of *D*, *RD*, or *AD*, or whose SQL routine body references *D*, *RD*, or *AD*.

- h) Let *R* be any SQL-invoked routine that is dependent on *D* and whose routine descriptor includes the descriptor of *D* or *RD*.
- i) *R* shall not be the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any of the following:
- 1) The SQL routine body of any routine descriptor.
  - 2) The <query expression> of any view descriptor.
  - 3) The <search condition> of any constraint descriptor.
  - 4) 04 The triggered action of any trigger descriptor.
- ii) The specific name of *R* shall not be included in any user-defined cast descriptor.
- iii) *R* shall not be the ordering function included in the descriptor of any user-defined type.

NOTE 335 — If CASCADE is specified, then such referenced objects will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.

NOTE 336 — The notion of an SQL-invoked routine being dependent on a user-defined type is defined in Subclause 4.27, “SQL-invoked routines”.

NOTE 337 — The notion of one data type being based on another data type is defined in Subclause 4.1, “Data types”.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of *D*.

## General Rules

- 1) Let *SN* be the <specific name> of any SQL-invoked routine that references *D*, *RD*, or *AD* or whose routine descriptor includes the descriptor of *D*, *RD*, or *AD* and that is not dependent on *D*. The following <drop routine statement> is effectively executed for each such SQL-invoked routine without further Access Rule checking:

DROP SPECIFIC ROUTINE *SN* CASCADE

NOTE 338 — The notion of an SQL-invoked routine being dependent on a user-defined type is defined in Subclause 4.27, “SQL-invoked routines”.

- 2) The following <drop transform statement> is effectively executed without further Access Rule checking:

DROP TRANSFORM ALL FOR *DN* CASCADE

NOTE 339 — This Rule should have no effect, since any external routine that depends on the transform being dropped also depends on the data type for which the transform is defined and hence should have already been dropped because of General Rule 1).

- 3) Let *UDCD* be the user-defined cast descriptor that references *DN* as the source data type. Let *TD* be the target data type included in *UDCD*. The following <drop user-defined cast statement> is effectively executed without further Access Rule checking:

DROP CAST ( *DN* AS *TD* ) CASCADE

**11.50 <drop data type statement>**

- 4) Let *UDCD* be the user-defined cast descriptor that references *DN* as the target data type. Let *SD* be the source data type included in *UDCD*. The following <drop user-defined cast statement> is effectively executed without further Access Rule checking:

```
DROP CAST (SD AS DN) CASCADE
```

- 5) Let *UDCD* be the user-defined cast descriptor that references the reference type whose referenced type is *DN* as the source data type. Let *TD* be the target data type included in *UDCD*. The following <drop user-defined cast statement> is effectively executed without further Access Rule checking:

```
DROP CAST (REF (DN) AS TD) CASCADE
```

- 6) Let *UDCD* be the user-defined cast descriptor that references the reference type whose referenced type is *DN* as the target data type. Let *SD* be the source data type included in *UDCD*. The following <drop user-defined cast statement> is effectively executed without further Access Rule checking:

```
DROP CAST (SD AS REF (DN)) CASCADE
```

- 7) For every privilege descriptor that references *D*, the following <revoke statement> is effectively executed:

```
REVOKE PRIV ON TYPE D FROM GRANTEE CASCADE
```

where *PRIV* and *GRANTEE* are respectively the action and grantee in the privilege descriptor.

- 8) The descriptor of every SQL-invoked routine that is said to be dependent on *D* is destroyed.

NOTE 340 — The notion of an SQL-invoked routine being dependent on a user-defined type is defined in [Subclause 4.27](#), “SQL-invoked routines”.

- 9) The descriptor of *D* is destroyed.

## Conformance Rules

- 1) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop data type statement> that contains a <drop behavior> that contains CASCADE.

## 11.51 <SQL-invoked routine>

*This Subclause is modified by Subclause 9.24, “<SQL-invoked routine>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 11.10, “<SQL-invoked routine>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 9.8, “<SQL-invoked routine>”, in ISO/IEC 9075-13.  
This Subclause is modified by Subclause 12.7, “<SQL-invoked routine>”, in ISO/IEC 9075-14.*

### Function

Define an SQL-invoked routine.

### Format

```
04 <SQL-invoked routine> ::=
 <schema routine>

<schema routine> ::=
 <schema procedure>
 | <schema function>

<schema procedure> ::=
 CREATE <SQL-invoked procedure>

<schema function> ::=
 CREATE <SQL-invoked function>

<SQL-invoked procedure> ::=
 PROCEDURE <schema qualified routine name> <SQL parameter declaration list>
 <routine characteristics>
 <routine body>

<SQL-invoked function> ::=
 { <function specification> | <method specification designator> } <routine body>

<SQL parameter declaration list> ::=
 <left paren> [<SQL parameter declaration>
 [{ <comma> <SQL parameter declaration> }...]] <right paren>

14 <SQL parameter declaration> ::=
 [<parameter mode>] [<SQL parameter name>] <parameter type> [RESULT]

<parameter mode> ::=
 IN
 | OUT
 | INOUT

14 <parameter type> ::=
 <data type> [<locator indication>]

<locator indication> ::=
 AS LOCATOR

<function specification> ::=
 FUNCTION <schema qualified routine name> <SQL parameter declaration list>
 <returns clause>
```

```

 <routine characteristics>
 [<dispatch clause>]

<method specification designator> ::=
 SPECIFIC METHOD <specific method name>
 | [INSTANCE | STATIC | CONSTRUCTOR]
 METHOD <method name> <SQL parameter declaration list>
 [<returns clause>]
 FOR <schema-resolved user-defined type name>

<routine characteristics> ::=
 [<routine characteristic>...]

<routine characteristic> ::=
 <language clause>
 | <parameter style clause>
 | SPECIFIC <specific name>
 | <deterministic characteristic>
 | <SQL-data access indication>
 | <null-call clause>
 | <returned result sets characteristic>
 | <savepoint level indication>

<savepoint level indication> ::=
 NEW SAVEPOINT LEVEL
 | OLD SAVEPOINT LEVEL

<returned result sets characteristic> ::=
 DYNAMIC RESULT SETS <maximum returned result sets>

<parameter style clause> ::=
 PARAMETER STYLE <parameter style>

<dispatch clause> ::=
 STATIC DISPATCH

14 <returns clause> ::=
 RETURNS <returns type>

<returns type> ::=
 <returns data type> [<result cast>]
 | <returns table type>

<returns table type> ::=
 TABLE <table function column list>

<table function column list> ::=
 <left paren> <table function column list element>
 [{ <comma> <table function column list element> }...] <right paren>

<table function column list element> ::=
 <column name> <data type>

<result cast> ::=
 CAST FROM <result cast from type>

<result cast from type> ::=
 <data type> [<locator indication>]

14 <returns data type> ::=

```

```

 <data type> [<locator indication>]

<routine body> ::=
 <SQL routine spec>
 | <external body reference>

<SQL routine spec> ::=
 [<rights clause>] <SQL routine body>

<rights clause> ::=
 SQL SECURITY INVOKER
 | SQL SECURITY DEFINER

<SQL routine body> ::=
 <SQL procedure statement>

<external body reference> ::=
 EXTERNAL [NAME <external routine name>]
 [<parameter style clause>]
 [<transform group specification>]
 [<external security clause>]

<external security clause> ::=
 EXTERNAL SECURITY DEFINER
 | EXTERNAL SECURITY INVOKER
 | EXTERNAL SECURITY IMPLEMENTATION DEFINED

13 <parameter style> ::=
 SQL
 | GENERAL

<deterministic characteristic> ::=
 DETERMINISTIC
 | NOT DETERMINISTIC

<SQL-data access indication> ::=
 NO SQL
 | CONTAINS SQL
 | READS SQL DATA
 | MODIFIES SQL DATA

<null-call clause> ::=
 RETURNS NULL ON NULL INPUT
 | CALLED ON NULL INPUT

<maximum returned result sets> ::=
 <unsigned integer>

<transform group specification> ::=
 TRANSFORM GROUP { <single group specification> | <multiple group specification> }

<single group specification> ::=
 <group name>

<multiple group specification> ::=
 <group specification> [{ <comma> <group specification> }...]

<group specification> ::=
 <group name> FOR TYPE <path-resolved user-defined type name>

```

## Syntax Rules

- 1) 09 An <SQL-invoked routine> specifies an *SQL-invoked routine*. Let *R* be the SQL-invoked routine specified by <SQL-invoked routine>.
- 2) If <SQL-invoked routine> immediately contains <schema routine>, then the SQL-invoked routine identified by <schema qualified routine name> is a *schema-level routine*.
- 3) 13 An <SQL-invoked routine> specified as an <SQL-invoked procedure> is called an *SQL-invoked procedure*; an <SQL-invoked routine> specified as an <SQL-invoked function> is called an *SQL-invoked function*. An <SQL-invoked function> that specifies a <method specification designator> is further called an *SQL-invoked method*. An SQL-invoked method that specifies STATIC is called a *static SQL-invoked method*. An SQL-invoked method that specifies CONSTRUCTOR is called an *SQL-invoked constructor method*.
- 4) If <returns type> *RST* specifies TABLE, then let *TCL* be the <table function column list> contained in <returns table type>.
  - a) For every <column name> *CN* contained in *TCL*, *CN* shall not be equivalent to any other <column name> contained in *TCL*.
  - b) *RST* is equivalent to the <returns type>

ROW *TCL* MULTISSET

- 5) If <SQL-invoked routine> specifies an SQL-invoked method, then
 

Case:

  - a) If a <specific method name> *SMN* is specified, then:
    - i) Case:
      - 1) If *SMN* does not contain <schema name>, then
 

Case:

        - A) If the <SQL-invoked routine> is contained in a <schema definition>, then the <schema name> that is specified or implicit in the <schema definition> is implicit.
        - B) Otherwise, the <schema name> that is specified or implicit for the <SQL-client module definition> is implicit.
      - 2) Otherwise, if <SQL-invoked routine> is contained in a <schema definition> then the <schema name> contained in *SMN* shall be equivalent to the specified or implicit <schema name> of the containing <schema definition>.
    - ii) Let *S* be the schema identified by the implicit or explicit <schema name> of *SMN*.
    - iii) There shall exist a method specification descriptor *DMS* included in the descriptor of a user-defined type *UDT* included in *S*, whose <specific method name> is *SMN*.
    - iv) Let *MN* be the number of SQL parameters in the unaugmented SQL parameter declaration list in *DMS*. *MN* is the number of SQL parameters in the unaugmented SQL parameter declaration list of *R*.
    - v) If *DMS* includes <result cast> *RC*, then *RC* is the <result cast> of *R*.



- vi) Let *SPN* be the <specific method name> in *DMS*. *SPN* is the <specific name> of *R*.
  - vii) Let *NPL* be the augmented SQL parameter declaration list of *DMS*. *NPL* is the augmented SQL parameter declaration list of *R*.
  - viii) Let *RN* be *SN*.<method name>, where *SN* is the <schema name> of the schema that includes the descriptor of *UDT*.
- b) Otherwise:
- i) 13 Let *UDTN* be the <schema-resolved user-defined type name> immediately contained in <method specification designator>. Let *UDT* be the user-defined type identified by *UDTN*.
  - ii) There shall exist a method specification descriptor *DMS* in the descriptor of *UDT* such that the <method name> of *DMS* is equivalent to the <method name>, *DMS* indicates STATIC if and only if the <method specification designator> specifies STATIC, *DMS* indicates CONSTRUCTOR if and only if the <method specification designator> specifies CONSTRUCTOR, and the declared type of every SQL parameter in the unaugmented SQL parameter declaration list in *DMS* is compatible with the declared type of the corresponding SQL parameter in the <SQL parameter declaration list> contained in the <method specification designator>. *DMS* identifies the corresponding method specification of the <method specification designator>.
  - iii) Let *MN* be the number of SQL parameters in the unaugmented SQL parameter declaration list in *DMS*.
  - iv) Let *PCOMS<sub>i</sub>*,  $1 \text{ (one)} \leq i \leq MN$ , be the *i*-th SQL parameter in the unaugmented SQL parameter declaration list of *DMS*. Let *POVMS<sub>i</sub>*,  $1 \text{ (one)} \leq i \leq MN$ , be the *i*-th SQL parameter contained in <method specification designator>.
  - v) For *i* varying from 1 (one) to *MN*, the <SQL parameter name>s contained in *PCOMS<sub>i</sub>* and *POVMS<sub>i</sub>* shall be equivalent.
  - vi) Let *PDMS<sub>i</sub>*,  $1 \text{ (one)} \leq i \leq MN$ , be the declared type of the *i*-th SQL parameter in the unaugmented SQL parameter declaration list in *DMS*. Let *PSM<sub>i</sub>* be the declared type of the *i*-th SQL parameter contained in <method specification designator>.
  - vii) With *i* ranging from 1 (one) to *MN*, the Syntax Rules of Subclause 9.18, “Data type identity”, are applied with *PDMS<sub>i</sub>* and *PSM<sub>i</sub>*.
  - viii) Case:
    - 1) If <returns clause> is specified, then let *RT* be the <returns data type> of *R*. Let *RDMS* be the <returns data type> in *DMS*. The Syntax Rules of Subclause 9.18, “Data type identity”, are applied with *RT* and *RDMS*.
    - 2) Otherwise, let *RDMS* be the <returns data type> of *R*.
  - ix) If *DMS* includes <result cast> *RC*, then
 

Case:

    - 1) If <returns clause> is specified, then <returns clause> shall contain <result cast>. Let *RDCT* be the <data type> specified in *RC*. Let *RCT* be the <data type> specified in the <result cast> contained in <returns clause>. The Syntax Rules of Subclause 9.18, “Data type identity”, are applied with *RDCT* and *RCT*.

- 2) Otherwise, *RC* is the <result cast> of *R*.
  - x) Let *SPN* be the <specific method name> in *DMS*. *SPN* is the <specific name> of *R*.
  - xi) Let *NPL* be the augmented SQL parameter declaration list of *DMS*.
  - xii) Let *RN* be *SN*.<method name>, where *SN* is the <schema name> of the schema that includes the descriptor of *UDT*.
- 6) If <SQL-invoked routine> specifies an SQL-invoked procedure or an SQL-invoked regular function, then:
- a) 13 <routine characteristics> shall contain at most one <language clause>, at most one <parameter style clause>, at most one <specific name>, at most one <deterministic characteristic>, at most one <SQL-data access indication>, at most one <null-call clause>, and at most one <returned result sets characteristic>.
  - b) <parameter style clause> shall not be specified both in <routine characteristics> and in <external body reference>.
  - c) The <routine characteristics> of a <function specification> shall not contain a <returned result sets characteristic>.
  - d) If the SQL-invoked routine is an SQL-invoked procedure and <returned result sets characteristic> is not specified, then DYNAMIC RESULT SETS 0 (zero) is implicit.
  - e) If <deterministic characteristic> is not specified, then NOT DETERMINISTIC is implicit.
  - f) Case:
    - i) If PROCEDURE is specified, then:
      - 1) <null-call clause> shall not be specified.
      - 2) <routine characteristics> shall not contain more than one <savepoint level indication>.
    - ii) Otherwise, if <null-call clause> is not specified, then CALLED ON NULL INPUT is implicit.
  - g) <SQL-data access indication> shall be specified.
  - h) 04 If <language clause> is not specified, then LANGUAGE SQL is implicit.
  - i) 13 An <SQL-invoked routine> that specifies or implies LANGUAGE SQL is called an *SQL routine*; an <SQL-invoked routine> that does not specify LANGUAGE SQL is called an *external routine*.
  - j) If <savepoint level indication> is specified, then PROCEDURE shall be specified.
  - k) If PROCEDURE is specified and <savepoint level indication> is not specified, then OLD SAVEPOINT LEVEL is implicit.
  - l) If NEW SAVEPOINT LEVEL is specified, then MODIFIES SQL DATA shall be specified.
  - m) If *R* is an SQL routine, then:
    - i) The <returns clause> shall not specify a <result cast>.
    - ii) <SQL-data access indication> shall not specify NO SQL.
    - iii) <parameter style clause> shall not be specified.

- n) An array-returning external function is an SQL-invoked function that is an external routine and that satisfies one of the following conditions:
  - i) A <result cast from type> is specified that simply contains an <array type> and does not contain a <locator indication>.
  - ii) A <result cast from type> is not specified and <returns data type> simply contains an <array type> and does not contain a <locator indication>.
- o) A multiset-returning external function is an SQL-invoked function that is an external routine and that satisfies one of the following conditions:
  - i) A <result cast from type> is specified that simply contains a <multiset type> and does not contain a <locator indication>.
  - ii) A <result cast from type> is not specified and <returns data type> simply contains a <multiset type> and does not contain a <locator indication>.
- p) Let *RN* be the <schema qualified routine name> of *R*.
- q) 04 If <SQL-invoked routine> is contained in a <schema definition> and *RN* contains a <schema name> *SN*, then *SN* shall be equivalent to the specified or implicit <schema name> of the containing <schema definition>. Let *S* be the SQL-schema identified by *SN*.
- r) Case:
  - i) If *R* is an SQL-invoked regular function and the <SQL parameter declaration list> contains an <SQL parameter declaration> that specifies a <data type> that is one of:
    - 1) A user-defined type.
    - 2) A collection type whose element type is a user-defined type.
    - 3) A collection type whose element type is a reference type.
    - 4) A reference type.then <dispatch clause> shall be specified.
  - ii) Otherwise, <dispatch clause> shall not be specified.
- s) If <specific name> is not specified, then an implementation-dependent <specific name> whose <schema name> is the equivalent to the <schema name> of *S* is implicit.
- t) If <specific name> contains a <schema name>, then that <schema name> shall be equivalent to the <schema name> of *S*. If <specific name> does not contain a <schema name>, then the <schema name> of *S* is implicit.
- u) The schema identified by the explicit or implicit <schema name> of the <specific name> shall not include a routine descriptor whose specific name is equivalent to <specific name> or a user-defined type descriptor that includes a method specification descriptor whose specific name is equivalent to <specific name>.
- v) If <returns data type> *RT* simply contains <locator indication>, then:
  - i) *R* shall be an external routine.

## 11.51 &lt;SQL-invoked routine&gt;

- ii) *RT* shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.
  - iii) <result cast> shall not be specified.
- w) If <result cast from type> *RCT* simply contains <locator indication>, then:
- i) *R* shall be an external routine.
  - ii) *RCT* shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.
- x) If *R* is an external routine, then:
- i) If <parameter style> is not specified, then PARAMETER STYLE SQL is implicit.
  - ii) 13 If *R* is an array-returning external function or a multiset-returning external function, then PARAMETER STYLE SQL shall be either specified or implied.
  - iii) 13 Case:
    - 1) If <transform group specification> is not specified, then a <multiple group specification> with a <group specification> *GS* for each <SQL parameter declaration> contained in <SQL parameter declaration list> whose <parameter type> *UDT* identifies a user-defined type with no <locator indication> is implicit. The <group name> of *GS* is implementation-defined and its <path-resolved user-defined type name> is *UDT*.
    - 2) If <single group specification> with a <group name> *GN* is specified, then <transform group specification> is equivalent to a <transform group specification> that contains a <multiple group specification> that contains a <group specification> *GS* for each <SQL parameter declaration> contained in <SQL parameter declaration list> whose <parameter type> *UDT* identifies a user-defined type with no <locator indication>. The <group name> of *GS* is *GN* and its <path-resolved user-defined type name> is *UDT*.
    - 3) Otherwise, <multiple group specification> is extended with a <group specification> *GS* for each <SQL parameter declaration> contained in <SQL parameter declaration list> whose <parameter type> *UDT* identifies a user-defined type with no <locator indication> and no equivalent of *UDT* is contained in any <group specification> contained in <multiple group specification>. The <group name> of *GS* is implementation-defined and its <path-resolved user-defined type name> is *UDT*.
  - iv) If a <result cast> is specified, then let *V* be some value of the <data type> specified in the <result cast> and let *RT* be the <returns data type>. The following shall be valid according to the Syntax Rules of Subclause 6.12, “<cast specification>”:

CAST ( *V* AS *RT* )

- y) 14 Let *NPL* be the <SQL parameter declaration list> contained in the <SQL-invoked routine>.
- 7) *NPL* specifies the list of SQL parameters of *R*. Each SQL parameter of *R* is specified by an <SQL parameter declaration>. If <SQL parameter name> is specified, then that SQL parameter of *R* is identified by an SQL parameter name.
- 8) *NPL* shall specify at most one <SQL parameter declaration> that specifies RESULT.

- 9) If  $R$  is an SQL-invoked function, then no <SQL parameter declaration> in  $NPL$  shall contain a <parameter mode>.
- 10) If  $R$  is an SQL routine, then every <SQL parameter declaration> in  $NPL$  shall contain an <SQL parameter name>.
- 11) No two <SQL parameter name>s contained in  $NPL$  shall be equivalent.
- 12) Let  $N$  and  $PN$  be the number of <SQL parameter declaration>s contained in  $NPL$ . For every <SQL parameter declaration>  $PD_i$ ,  $1 \text{ (one)} \leq i \leq N$ :
  - a) <parameter type>  $PT_i$  immediately contained in  $PD_i$  shall not specify ROW.
  - b) If  $PT_i$  simply contains <locator indication>, then:
    - i)  $R$  shall be an external routine.
    - ii)  $PT_i$  shall specify either binary large object type, character large object type, array type, multiset type, or user-defined type.
  - c) If  $PD_i$  immediately contains RESULT, then:
    - i)  $R$  shall be an SQL-invoked function.
    - ii)  $PT_i$  shall specify a structured type  $ST$ . Let  $STN$  be the <user-defined type name> that identifies  $ST$ .
    - iii) The <returns data type> shall specify  $STN$ .
    - iv)  $R$  is a type-preserving function and  $PD_i$  specifies the result SQL parameter of  $R$ .
  - d) If  $PD_i$  does not contain a <parameter mode>, then a <parameter mode> that specifies IN is implicit.
  - e) Let  $P_i$  be the  $i$ -th SQL parameter.

Case:

    - i) If the <parameter mode> specifies IN, then  $P_i$  is an input SQL parameter.
    - ii) If the <parameter mode> specifies OUT, then  $P_i$  is an output SQL parameter.
    - iii) If the <parameter mode> specifies INOUT, then  $P_i$  is both an input SQL parameter and an output SQL parameter.
- 13) The scope of  $RN$  is the <routine body> of  $R$ .
- 14) The scope of an <SQL parameter name> contained in  $NPL$  is the <routine body>  $RB$  of the <SQL-invoked procedure> or <SQL-invoked function> that contains  $NPL$ .
- 15) An <SQL-invoked routine> shall not contain a <host parameter name>, a <dynamic parameter specification>, or an <embedded variable name>.
- 16) Case:
  - a) If  $R$  is an SQL-invoked procedure, then  $S$  shall not include another SQL-invoked procedure whose <schema qualified routine name> is equivalent to  $RN$  and whose number of SQL parameters is  $PN$ .

- b) Otherwise:
  - i) Case:
    - 1) If  $R$  is a static SQL-invoked method, then let  $SCR$  be the set containing every static SQL-invoked method of type  $UDT$ , including  $R$ , whose <schema qualified routine name> is equivalent to  $RN$  and whose number of SQL parameters is  $PN$ .
    - 2) If  $R$  is an SQL-invoked constructor method, then let  $SCR$  be the set containing every SQL-invoked constructor method of type  $UDT$ , including  $R$ , whose <schema qualified routine name> is equivalent to  $RN$  and whose number of SQL parameters is  $PN$ .
    - 3) Otherwise, let  $SCR$  be the set containing every SQL-invoked function in  $S$  that is neither a static SQL-invoked method nor an SQL-invoked constructor method, including  $R$ , whose <schema qualified routine name> is equivalent to  $RN$  and whose number of SQL parameters is  $PN$ .
  - ii) Let  $AL$  be an <SQL argument list> constructed from a list of arbitrarily-selected values in which the declared type of every value  $A_i$  in  $AL$  is compatible with the declared type of the corresponding SQL parameter  $P_i$  of  $R$ .
  - iii) For every  $A_i$ , eliminate from  $SCR$  every SQL-invoked routine  $SIR$  for which the type designator of the declared type of the SQL parameter  $P_i$  of  $SIR$  is not in the type precedence list of the declared type of  $A_i$ .
  - iv) The Syntax Rules of Subclause 9.4, “Subject routine determination”, are applied with <SQL argument list> as *SQL ARGUMENT LIST* and the set of SQL-invoked routines as *SET OF SQL-INVOKED ROUTINES*, yielding  $SR$  returned from *SET OF SUBJECT ROUTINES*.
- 17) If  $R$  is an SQL-invoked method but not a static SQL-invoked method, then the first SQL parameter of  $NPL$  is called the *subject parameter* of  $R$ .
- 18) If  $R$  is an SQL-invoked regular function  $F$  whose first SQL parameter has a declared type that is a user-defined type, then:
  - a) Let  $UDT$  be the declared type of the first SQL parameter of  $F$ .
  - b) Let  $DMS$  be a method specification descriptor of an instance method in the descriptor of  $UDT$  such that:
    - i) The <schema qualified routine name> of  $F$  and the <routine name> of  $DMS$  have equivalent <qualified identifier>s.
    - ii)  $F$  and the augmented SQL parameter declaration list of  $DMS$  have the same number of SQL parameters.
  - c) Let  $PDMS_i$ ,  $1 \text{ (one)} \leq i \leq PN$ , be the declared type of the  $i$ -th SQL parameter in the unaugmented SQL parameter declaration list in  $DMS$  and let  $PMS_i$  be the declared type of the  $i$ -th SQL parameter contained in <function specification>.
  - d) One of the following conditions shall be false:
    - i) The declared type of  $PDMS_i$ ,  $1 \text{ (one)} \leq i \leq N$  is compatible with the declared type of SQL parameter  $PMS_{i+1}$ .

ii) *UDT* is a subtype or a supertype of the declared type of *PMS*<sub>1</sub>.

19) If *R* is an SQL routine, then:

- a) <SQL routine spec> shall be specified.
- b) If <rights clause> is not specified, then SQL SECURITY DEFINER is implicit.
- c) If READS SQL DATA is specified, then it is implementation-defined whether the <SQL routine body> shall not contain an <SQL procedure statement> *S* that satisfies at least one of the following:
  - i) *S* is an <SQL data change statement>.
  - ii) *S* contains a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.
  - iii) *S* contains an <SQL procedure statement> that is an <SQL data change statement>.
- d) If CONTAINS SQL is specified, then it is implementation-defined whether the <SQL routine body> shall not contain an <SQL procedure statement> *S* that satisfies at least one of the following:
  - i) *S* is an <SQL data statement> other than <free locator statement> and <hold locator statement>.
  - ii) *S* contains a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data or possibly reads SQL-data.
  - iii) *S* contains an <SQL procedure statement> that is an <SQL data statement> other than <free locator statement> and <hold locator statement>.
- e) If DETERMINISTIC is specified, then it is implementation-defined whether the <SQL routine body> shall not contain an <SQL procedure statement> that is possibly non-deterministic.
- f) It is implementation-defined whether the <SQL routine body> shall not contain an <SQL connection statement>, an <SQL schema statement>, an <SQL dynamic statement>, or an <SQL transaction statement> other than a <savepoint statement>, <release savepoint statement>, or a <rollback statement> that specifies a <savepoint clause>.

NOTE 341 — Conforming SQL language shall not contain an <SQL connection statement> or an <SQL transaction statement> other than a <savepoint statement>, a <release savepoint statement>, or a <rollback statement> that specifies a <savepoint clause>, but an implementation is not required to treat this as a syntax error.

- g) An <SQL routine body> shall not immediately contain an <SQL procedure statement> that simply contains a <schema definition>.
- h) An <SQL routine body> shall not immediately contain an <SQL procedure statement> that simply contains a <table reference> that identifies a declared local temporary table.

20) If *R* is an external routine, then:

- a) <SQL routine spec> shall not be specified.
- b) If <external security clause> is not specified, then EXTERNAL SECURITY IMPLEMENTATION DEFINED is implicit.
- c) If an <external routine name> is not specified, then an <external routine name> that is equivalent to the <qualified identifier> of *R* is implicit.
- d) If PARAMETER STYLE SQL is specified, then:

- i) Case:
  - 1) If  $R$  is an array-returning external function or a multiset-returning external function with the element type being a row type, then let  $FRN$  be the degree of the element type.
  - 2) Otherwise, let  $FRN$  be 1 (one).
- ii) If  $R$  is an array-returning external function or a multiset-returning external function, then let  $AREF$  be  $FRN+6$ . Otherwise, let  $AREF$  be  $FRN+4$ .
- iii) If  $R$  is an SQL-invoked function, then let the *effective SQL parameter list* be a list of  $PN+FRN+N+AREF$  SQL parameters, as follows:

- 1) For  $i$  ranging from 1 (one) to  $PN$ , the  $i$ -th effective SQL parameter list entry is defined as follows.

Case:

- A) 14 If the <parameter type>  $T_i$  simply contained in the  $i$ -th <SQL parameter declaration> contains <locator indication>, then the  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by INTEGER.
- B) If the <parameter type>  $T_i$  immediately contained in the  $i$ -th <SQL parameter declaration> is a <path-resolved user-defined type name> without a <locator indication>, then:

I) Case:

- 1) If  $R$  is an SQL-invoked method that is an overriding method, then the Syntax Rules of Subclause 9.20, “Determination of a from-sql function for an overriding method”, are applied with  $R$  and  $i$  as *ROUTINE* and *POSITION*, respectively. There shall be an applicable from-sql function  $FSF_i$ .
- 2) Otherwise, the Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, are applied with the data type identified by  $T_i$ , and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSF_i$ .

II)  $FSF_i$  is called the *from-sql function associated with the  $i$ -th SQL parameter*.

III) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by the <returns data type> of  $FSF_i$ .

- C) Otherwise, the  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration>.

2) Case:

- A) If  $FRN$  is 1 (one), then effective SQL parameter list entry  $PN+FRN$  has <parameter mode> OUT; its <parameter type>  $PT$  is defined as follows:
  - I) If <result cast> is specified, then let  $RT$  be <result cast from type>; otherwise, let  $RT$  be <returns data type>.
  - II) Case:
    - 1) 14 If  $RT$  simply contains <locator indication>, then  $PT$  is INTEGER.



- 2) If *RT* specifies a <path-resolved user-defined type name> without a <locator indication>, then:
    - a) Case:
      - i) If *R* is an SQL-invoked method that is an overriding method, then the Syntax Rules of Subclause 9.22, “Determination of a to-sql function for an overriding method”, are applied with *R* as *ROUTINE*. There shall be an applicable to-sql function *TSF*.
      - ii) Otherwise, the Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with the data type identified by *RT* and the <group name> contained in the <group specification> that contains *RT* as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function *TSF*.
    - b) *TSF* is called the *to-sql function* associated with the result.
    - c) Case:
      - i) If *TSF* is an SQL-invoked method, then *PT* is the <parameter type> of the second SQL parameter of *TSF*.
      - ii) Otherwise, *PT* is the <parameter type> of the first SQL parameter of *TSF*.
  - 3) If *R* is an array-returning external function or a multiset-returning external function, then let *PT* be the element type of *RT*.
  - 4) Otherwise, *PT* is *RT*.
- B) Otherwise, for *i* ranging from *PN*+1 to *PN*+*FRN*, the *i*-th effective SQL parameter list entry is defined as follows.
- Case:
- I) Its <parameter mode> is OUT.
  - II) Let *RFT<sub>i-PN</sub>* be the data type of the *i-PN*-th field of the element type of the <returns data type>. The <parameter type> *PT<sub>i</sub>* of the *i*-th effective SQL parameter list entry is determined as follows:
    - 1) If *RFT<sub>i-PN</sub>* specifies a <path-resolved user-defined type name>, then:
      - a) Case:
        - i) If *R* is an SQL-invoked method that is an overriding method, then the Syntax Rules of Subclause 9.22, “Determination of a to-sql function for an overriding method”, are applied with *R* as *ROUTINE*. There shall be an applicable to-sql function *TSF*.
        - ii) Otherwise, the Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with the data type identified by *RFT<sub>i-PN</sub>* and the <group name> contained in the <group specification> that contains *RFT<sub>i-PN</sub>* as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function *TSF*.

- b) *TSF* is called the *to-sql function* associated with  $RFT_{i-PN}$ .
- c) Case:
  - i) If *TSF* is an SQL-invoked method, then  $PT_i$  is the <parameter type> of the second SQL parameter of *TSF*.
  - ii) Otherwise,  $PT_i$  is the <parameter type> of the first SQL parameter of *TSF*.

2) 14 Otherwise,  $PT_i$  is  $RFT_{i-PN}$ .

- 3) Effective SQL parameter list entries  $(PN+FRN)+1$  to  $(PN+FRN)+N+FRN$  are  $N+FRN$  occurrences of SQL parameters of an implementation-defined <data type> that is an exact numeric type with scale 0 (zero). For  $i$  ranging from  $(PN+FRN)+1$  to  $(PN+FRN)+N+FRN$ , the <parameter mode> for the  $i$ -th such effective SQL parameter is the same as that of the  $i-FRN-PN$ -th effective SQL parameter.
- 4) Effective SQL parameter list entry  $(PN+FRN)+(N+FRN)+1$  is an SQL parameter of a <data type> that is character string of length 5 and the character set specified for SQLSTATE values, with <parameter mode> INOUT.

NOTE 342 — The character set specified for SQLSTATE values is defined in Subclause 24.1, “SQLSTATE”.

- 5) Effective SQL parameter list entry  $(PN+FRN)+(N+FRN)+2$  is an SQL parameter of a <data type> that is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> IN.
- 6) Effective SQL parameter list entry  $(PN+FRN)+(N+FRN)+3$  is an SQL parameter of a <data type> that is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> IN.
- 7) Effective SQL parameter list entry  $(PN+FRN)+(N+FRN)+4$  is an SQL parameter of a <data type> that is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> INOUT.
- 8) If  $R$  is an array-returning external function or a multiset-returning external function, then:
  - A) Effective SQL parameter type list entry  $(PN+FRN)+(N+FRN)+5$  is an SQL parameter whose <data type> is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> INOUT.
  - B) Effective SQL parameter type list entry  $(PN+FRN)+(N+FRN)+6$  is an SQL parameter whose <data type> is an exact numeric type with scale 0 (zero) and with <parameter mode> IN.
- iv) If  $R$  is an SQL-invoked procedure, then let the *effective SQL parameter list* be a list of  $PN+N+4$  SQL parameters, as follows:
  - 1) For  $i$  ranging from 1 (one) to  $PN$ , the  $i$ -th effective SQL parameter list entry is defined as follows.
    - Case:
      - A) 14 If the <parameter type>  $T_i$  simply contained in the  $i$ -th <SQL parameter declaration> simply contains <locator indication>, then the  $i$ -th effective SQL parameter list entry

is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by INTEGER.

- B) If the <parameter type>  $T_i$  simply contained in the  $i$ -th <SQL parameter declaration> is a <path-resolved user-defined type name> without a <locator indication>, then:
- I) Case:
- 1) If the <parameter mode> immediately contained in the  $i$ -th <SQL parameter declaration> is IN, then:
    - a) The Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSF_i$ .  $FSF_i$  is called the *from-sql function* associated with the  $i$ -th SQL parameter.
    - b) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by the <returns data type> of  $FSF_i$ .
  - 2) If the <parameter mode> immediately contained in the  $i$ -th <SQL parameter declaration> is OUT, then:
    - a) The Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function  $TSF_i$ .  $TSF_i$  is called the *to-sql function* associated with  $i$ -th SQL parameter.
    - b) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by  
Case:
      - i) If  $TSF_i$  is an SQL-invoked method, then the <parameter type> of the second SQL parameter of  $TSF_i$ .
      - ii) Otherwise, the <parameter type> of the first SQL parameter of  $TSF_i$ .
  - 3) Otherwise:
    - a) The Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSF_i$ .  $FSF_i$  is called the *from-sql function* associated with the  $i$ -th SQL parameter.
    - b) The Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with the data type identified by  $T_i$  and the <group name>

contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function  $TSF_i$ .  $TSF_i$  is called the *to-sql function associated with the  $i$ -th SQL parameter*.

- c) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by the <returns data type> of  $FSF_i$ .
- C) Otherwise, the  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration>.
- 2) Effective SQL parameter list entries  $PN+1$  to  $PN+N$  are  $N$  occurrences of an SQL parameter of an implementation-defined <data type> that is an exact numeric type with scale 0. The <parameter mode> for the  $i$ -th such effective SQL parameter is the same as that of the  $i-PN$ -th effective SQL parameter.
- 3) Effective SQL parameter list entry  $(PN+N)+1$  is an SQL parameter of a <data type> that is character string of length 5 and character set SQL\_TEXT with <parameter mode> INOUT.
- 4) Effective SQL parameter list entry  $(PN+N)+2$  is an SQL parameter of a <data type> that is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> IN.
- 5) Effective SQL parameter list entry  $(PN+N)+3$  is an SQL parameter of a <data type> that is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> IN.
- 6) Effective SQL parameter list entry  $(PN+N)+4$  is an SQL parameter of a <data type> that is character string of implementation-defined length and character set SQL\_TEXT with <parameter mode> INOUT.
- e) 13 If PARAMETER STYLE GENERAL is specified, then let the *effective SQL parameter list* be a list of  $PN$  parameters such that, for  $i$  ranging from 1 (one) to  $PN$ , the  $i$ -th effective SQL parameter list entry is defined as follows.

Case:

- i) 14 If the <parameter type>  $T_i$  simply contained in the  $i$ -th <SQL parameter declaration> simply contains <locator indication>, then the  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by INTEGER.
- ii) If the <parameter type>  $T_i$  simply contained in the  $i$ -th <SQL parameter declaration> is a <path-resolved user-defined type name> without a <locator indication>, then:
  - 1) Case:
    - A) If the <parameter mode> immediately contained in the  $i$ -th <SQL parameter declaration> is IN, then:
      - I) The Syntax Rules of [Subclause 9.19](#), “Determination of a from-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSF_i$ .  $FSF_i$  is called the *from-sql function associated with the  $i$ -th SQL parameter*.

- II) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by the <returns data type> of  $FSF_i$ .
  - B) If the <parameter mode> immediately contained in the  $i$ -th <SQL parameter declaration> is OUT, then:
    - I) The Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function  $TSF_i$ .  $TSF_i$  is called the *to-sql function associated with the  $i$ -th SQL parameter*.
    - II) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by
 

Case:

      - 1) If  $TSF_i$  is an SQL-invoked method, then the <parameter type> of the second SQL parameter of  $TSF_i$ .
      - 2) Otherwise, the <parameter type> of the first SQL parameter of  $TSF_i$ .
  - C) Otherwise:
    - I) The Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSF_i$ .  $FSF_i$  is called the *from-sql function associated with the  $i$ -th SQL parameter*.
    - II) The Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with the data type identified by  $T_i$  and the <group name> contained in the <group specification> that contains  $T_i$  as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function  $TSF_i$ .  $TSF_i$  is called the *to-sql function associated with the  $i$ -th SQL parameter*.
    - III) The  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration> with the <parameter type> replaced by the <returns data type> of  $FSF_i$ .
  - iii) Otherwise, the  $i$ -th effective SQL parameter list entry is the  $i$ -th <SQL parameter declaration>.
- NOTE 343 — If the SQL-invoked routine is an SQL-invoked function, then the value returned from the external routine is passed to the SQL-implementation in an implementation-dependent manner. An SQL parameter is not used for this purpose.
- f) Depending on whether the <language clause> specifies ADA, C, COBOL, FORTRAN, M, PASCAL, or PLI, let the *operative data type correspondences table* be Table 16, “Data type correspondences for Ada”, Table 17, “Data type correspondences for C”, Table 18, “Data type correspondences for COBOL”, Table 19, “Data type correspondences for Fortran”, Table 20, “Data type correspondences for M”, Table 21, “Data type correspondences for Pascal”, or Table 22, “Data type correspondences for PL/I”, respectively. Refer to the two columns of the operative data type correspondences table as the “SQL data type” column and the “host data type column”.

- g) 13 Any <data type> in an effective SQL parameter list entry shall specify a data type listed in the SQL data type column for which the corresponding row in the host data type column is not “None”.

21) 13 Case:

- a) If <method specification designator> is specified, then:
- i) *R* is deterministic if *DMS* indicates that the method is deterministic; otherwise, *R* is possibly non-deterministic.
  - ii) *R* possibly modifies SQL-data if the SQL-data access indication of *DMS* indicates that the method possibly modifies SQL-data. *R* possibly reads SQL-data if the SQL-data access indication of *DMS* indicates that the method possibly reads SQL-data. *R* possibly contains SQL if the SQL-data access indication of *DMS* indicates that the method possibly contains SQL. Otherwise, *R* does not possibly contain SQL.
- b) Otherwise:
- i) If DETERMINISTIC is specified, then *R* is *deterministic*; otherwise, it is *possibly non-deterministic*.
  - ii) An <SQL-invoked routine> *possibly modifies SQL-data* if and only if <SQL-data access indication> specifies MODIFIES SQL DATA.
  - iii) An <SQL-invoked routine> *possibly reads SQL-data* if and only if <SQL-data access indication> specifies READS SQL DATA.
  - iv) An <SQL-invoked routine> *possibly contains SQL* if and only if <SQL-data access indication> specifies CONTAINS SQL.
  - v) An <SQL-invoked routine> *does not possibly contain SQL* if and only if <SQL-data access indication> specifies NO SQL.

22) If *R* is a schema-level routine, then let the containing schema be the schema identified by the <schema name> explicitly or implicitly contained in <schema qualified routine name>.

23) If the <SQL-invoked routine> is contained in a <schema definition>, then let *A* be the explicit or implicit <authorization identifier> of the <schema definition>; otherwise, let *A* be the <authorization identifier> that owns the schema identified by the explicit or implicit <schema name> of the <schema qualified routine name>.

## Access Rules

- 1) 13 If an <SQL-invoked routine> is contained in an <SQL-client module definition> *M* with no intervening <schema definition>, then the enabled authorization identifiers shall include the <authorization identifier> that owns *S*.
- 2) If *R* is an external routine and if any of its SQL parameters have an associated from-sql function or a to-sql function, or if *R* has a to-sql function associated with the result, then

Case:

- a) If <SQL-invoked routine> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges of the <authorization identifier> that owns the containing schema shall include EXECUTE on all from-sql

functions (if any) and on all to-sql functions (if any) associated with the SQL parameters and on the to-sql function associated with the result (if any).

- b) Otherwise, the current privileges shall include EXECUTE on all from-sql functions (if any) and on all to-sql functions (if any) associated with the SQL parameters and on the to-sql function associated with the result (if any).

## General Rules

- 1) If *R* is a schema-level routine, then a privilege descriptor is created that defines the EXECUTE privilege on *R* to the <authorization identifier> that owns the schema that includes *R*. The grantor for the privilege descriptor is set to the special grantor value “\_SYSTEM”. This privilege is grantable if and only if one of the following is satisfied:
  - a) *R* is an SQL routine and all of the privileges necessary for the <authorization identifier> to successfully execute the <SQL procedure statement> contained in the <routine body> are grantable. The necessary privileges include the EXECUTE privilege on every subject routine of every <routine invocation> contained in the <SQL procedure statement>.
  - b) *R* is an SQL routine and SQL SECURITY INVOKER is specified.
  - c) *R* is an external routine.
- 2) Case:
  - a) If <SQL-invoked routine> is contained in a <schema definition>, then let *DP* be the SQL-path of that <schema definition>.
  - b) If <SQL-invoked routine> is contained in a <preparable statement> or in a <direct SQL statement>, then let *DP* be the SQL-path of the current SQL-session.
  - c) Otherwise, let *DP* be the SQL-path of the <SQL-client module definition> that contains <SQL-invoked routine>.
- 3) If <method specification designator> is not specified, then a routine descriptor is created that describes the SQL-invoked routine being defined:
  - a) The routine name included in the routine descriptor is <schema qualified routine name>.
  - b) The specific name included in the routine descriptor is <specific name>.
  - c) The routine descriptor includes, for each SQL parameter in *NPL*, the name, declared type, ordinal position, an indication of whether the SQL parameter is input, output, or both, and an indication of whether the SQL parameter is a result SQL parameter.
  - d) The maximum number of returned result sets included in the routine descriptor is  
Case:
    - i) If the SQL-invoked routine is an SQL-invoked procedure, then the explicit or implicit value of <maximum returned result sets>.
    - ii) Otherwise, 0 (zero).
  - e) The routine descriptor includes an indication of whether the SQL-invoked routine is an SQL-invoked function or an SQL-invoked procedure.



- f) If the SQL-invoked routine is an SQL-invoked function, then:
  - i) The routine descriptor includes an indication that the SQL-invoked function is not an SQL-invoked method.
  - ii) The routine descriptor includes the data type in the <returns data type>. If the <returns data type> simply contains <locator indication>, then the routine descriptor includes an indication that the return value is a locator.
  - iii) The SQL-invoked routine descriptor includes an indication of whether the SQL-invoked routine is a null-call function.
- g) If the SQL-invoked routine is a type-preserving function, then the routine descriptor includes an indication that the SQL-invoked routine is a type-preserving function.
- h) The name of the language in which the body of the SQL-invoked routine was written is the <language name> contained in the <language clause>.
- i) If the SQL-invoked routine is an SQL routine, then the SQL routine body of the routine descriptor is the <SQL routine body>.
- j) If the SQL-invoked routine is an SQL-invoked function or NEW SAVEPOINT LEVEL is specified, then the routine descriptor includes an indication that a new savepoint level is to be established whenever the routine is invoked.

NOTE 344 — The use of savepoint levels is dependent on Feature T272, “Enhanced savepoint management”.

- k) Case:
  - i) If SQL SECURITY INVOKER is specified, then the SQL security characteristic in the routine descriptor is INVOKER.
  - ii) Otherwise, the SQL security characteristic in the routine descriptor is DEFINER.
- l) If the SQL-invoked routine is an external routine, then:
  - i) The external name of the routine descriptor is <external routine name>.
  - ii) 13 The routine descriptor includes an indication of whether the *parameter passing style* is PARAMETER STYLE SQL or PARAMETER STYLE GENERAL.
- m) The SQL-invoked routine descriptor includes an indication of whether the SQL-invoked routine is DETERMINISTIC or NOT DETERMINISTIC.
- n) The routine descriptor includes an indication of whether the SQL-invoked routine does not possibly contain SQL, possibly contains SQL, possibly reads SQL-data, or possibly modifies SQL-data.
- o) If the SQL-invoked routine specifies a <result cast>, then the routine descriptor includes an indication that the SQL-invoked routine specifies a <result cast> and the <data type> specified in the <result cast>. If <result cast> contains <locator indication>, then the routine descriptor includes an indication that the <data type> specified in the <result cast> has a locator indication.
- p) For every SQL parameter that has an associated from-sql function *FSF*, the routine descriptor includes the specific name of *FSF*.
- q) 14 For every SQL parameter that has an associated to-sql function *TSF*, the routine descriptor includes the specific name of *TSF*.



- r) 14 If *R* is an external function and if *R* has a to-sql function associated with its result *TRF*, then the routine descriptor includes the specific name of *TRF*.
- s) For every SQL parameter whose <SQL parameter declaration> contains <locator indication>, the routine descriptor includes an indication that the SQL parameter is a locator parameter.
- t) The routine authorization identifier included in the routine descriptor is the <authorization identifier> that owns *S*.
- u) The routine SQL-path included in the routine descriptor is *DP*.

NOTE 345 — The routine SQL-path is used to set the routine SQL-path of the current SQL-session when *R* is invoked. The routine SQL-path of the current SQL-session is used by the Syntax Rules of Subclause 10.4, “<routine invocation>”, to define the subject routines of <routine invocation>s contained in *R*. The same routine SQL-path is used whenever *R* is invoked.

- v) 04 The routine descriptor includes an indication that the routine is a schema-level routine.
- w) The routine descriptor includes an indication of whether the SQL-invoked routine is dependent on a user-defined type.

NOTE 346 — The notion of an SQL-invoked routine being dependent on a user-defined type is defined in Subclause 4.27, “SQL-invoked routines”.

- 4) If <method specification designator> is specified, then let *DMS* be the descriptor of the corresponding method specification. A routine descriptor is created that describes the SQL-invoked routine being defined.
  - a) The routine name included in the routine descriptor is *RN*.
  - b) The specific name included in the routine descriptor is <specific name>.
  - c) The routine descriptor includes, for each SQL parameter in *NPL*, the name, declared type, ordinal position, an indication of whether the SQL parameter is input, output, or both, and an indication of whether the SQL parameter is a result SQL parameter.
  - d) The routine descriptor includes an indication that the SQL-invoked routine is an SQL-invoked function that is an SQL-invoked method, the name of the user-defined type *UDT*, and an indication of whether *STATIC* or *CONSTRUCTOR* was specified.
  - e) If the SQL-invoked routine is a type-preserving function, then the routine descriptor includes an indication that the SQL-invoked routine is a type-preserving function.
  - f) If the SQL-invoked routine is a mutator function, then the routine descriptor includes an indication that the SQL-invoked routine is a mutator function.
  - g) The routine descriptor includes the data type in the <returns data type>.
  - h) The name of the language in which the body of the SQL-invoked routine was written is the <language name> contained in the <language clause> in *DMS*.
  - i) If the SQL-invoked routine is an SQL routine, then the SQL routine body of the routine descriptor is the <SQL routine body>.
  - j) Case:
    - i) If *SQL SECURITY INVOKER* is specified, then the SQL security characteristic in the routine descriptor is *INVOKER*.
    - ii) Otherwise, the SQL security characteristic in the routine descriptor is *DEFINER*.

- k) If the SQL-invoked routine is an external routine, then:
  - i) The external name of the routine descriptor is <external routine name>.
  - ii) The routine descriptor includes an indication of whether the parameter passing style is **PARAMETER STYLE SQL** or **PARAMETER STYLE GENERAL**, which is the same as the indication of <parameter style> in *DMS*.
- l) The routine descriptor includes an indication of whether the SQL-invoked routine is deterministic.
- m) The routine descriptor includes an indication of whether the SQL-invoked routine possibly modifies SQL-data, possibly read SQL-data, possibly contains SQL, or does not possibly contain SQL.
- n) The routine descriptor includes an indication of whether the SQL-invoked routine is a null-call function, which is the same as the indication in *DMS*.
- o) If *DMS* specifies a <result cast>, then the routine descriptor includes an indication that the SQL-invoked routine specifies a <result cast> and the <data type> specified in the <result cast> of *DMS*.
- p) The routine authorization identifier included in the routine descriptor is the <authorization identifier> that owns *S*.
- q) The routine SQL-path included in the routine descriptor is *DP*.

NOTE 347 — The routine SQL-path is used to set the routine SQL-path of the current SQL-session when *R* is invoked. The routine SQL-path of the current SQL-session is used by the Syntax Rules of Subclause 10.4, “<routine invocation>”, to define the subject routine of <routine invocation>s contained in *R*. The same routine SQL-path is used whenever *R* is invoked.

- r) The routine descriptor includes an indication of whether the routine is a schema-level routine.
- s) The routine descriptor includes an indication of whether the SQL-invoked routine is dependent on a user-defined type.

NOTE 348 — The notion of an SQL-invoked routine being dependent on a user-defined type is defined in Subclause 4.27, “SQL-invoked routines”.

- 5) The creation timestamp and the last-altered timestamp included in the routine descriptor are the values of **CURRENT\_TIMESTAMP**.
- 6) If *R* is an external routine, then the routine descriptor of *R* includes further elements determined as follows:
  - a) Case:

- i) 13 If <SQL-data access indication> in the descriptor of *R* is **MODIFIES SQL DATA**, **READS SQL DATA**, or **CONTAINS SQL**, then:
  - 1) Let *P* be the program identified by the <external routine name>.
  - 2) The external routine SQL-path is the <schema name list> immediately contained in the <path specification> that is immediately contained in the <module path specification> of the <SQL-client module definition> of *P*.
- ii) Otherwise, the external routine SQL-path is implementation-defined.

NOTE 349 — The external routine SQL-path is used to set the routine SQL-path of the current SQL-session when *R* is invoked. The routine SQL-path of the current SQL-session is used by the Syntax Rules of Subclause 10.4, “<routine invocation>”, to define the subject routines of <routine invocation>s contained in the <SQL-client module definition> of *P*. The same external routine SQL-path is used whenever *R* is invoked.

- b) The external security characteristic in the routine descriptor is

Case:

- i) If <external security clause> specifies EXTERNAL SECURITY DEFINER, then DEFINER.
  - ii) If <external security clause> specifies EXTERNAL SECURITY INVOKER, then INVOKER.
  - iii) Otherwise, EXTERNAL SECURITY IMPLEMENTATION DEFINED.
- c) The effective SQL parameter list is the *effective SQL parameter list*.

## Conformance Rules

- 1) Without Feature T471, “Result sets return value”, conforming SQL language shall not contain a <returned result sets characteristic>.
- 2) Without Feature T322, “Overloading of SQL-invoked functions and procedures”, conforming SQL language shall not contain a <schema routine> in which the schema identified by the explicit or implicit schema name of the <schema qualified routine name> includes a routine descriptor whose routine name is <schema qualified routine name>.
- 3) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <method specification designator>.
- 4) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <transform group specification>.
- 5) Without Feature S024, “Enhanced structured types”, an <SQL parameter declaration> shall not contain RESULT.
- 6) Without Feature T571, “Array-returning external SQL-invoked functions”, conforming SQL language shall not contain an <SQL-invoked routine> that defines an array-returning external function.
- 7) Without Feature T572, “Multiset-returning external SQL-invoked functions”, conforming SQL language shall not contain an <SQL-invoked routine> that defines a multiset-returning external function.
- 8) Without Feature S201, “SQL-invoked routines on arrays”, conforming SQL language shall not contain a <parameter type> that is based on an array type.
- 9) Without Feature S201, “SQL-invoked routines on arrays”, conforming SQL language shall not contain a <returns data type> that is based on an array type.
- 10) Without Feature S202, “SQL-invoked routines on multisets”, conforming SQL language shall not contain a <parameter type> that is based on a multiset type.
- 11) Without Feature S202, “SQL-invoked routines on multisets”, conforming SQL language shall not contain a <returns data type> that is based on a multiset type.
- 12) Without Feature T323, “Explicit security for external routines”, conforming SQL language shall not contain an <external security clause>.
- 13) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies a structured type.

- 14) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies a structured type.
- 15) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies an array type.
- 16) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies an array type.
- 17) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies a multiset type.
- 18) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies a multiset type.
- 19) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies a large object type.
- 20) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies a large object type.
- 21) Without Feature S027, “Create method by specific method name”, conforming SQL language shall not contain a <method specification designator> that contains SPECIFIC METHOD.
- 22) Without Feature T324, “Explicit security for SQL routines”, conforming SQL language shall not contain a <rights clause>.
- 23) Without Feature T326, “Table functions”, conforming SQL language shall not contain a <returns table type>.
- 24) Without Feature T651, “SQL-schema statements in SQL routines”, conforming SQL language shall not contain an <SQL routine body> that contains an SQL-schema statement.
- 25) Without Feature T652, “SQL-dynamic statements in SQL routines”, conforming SQL language shall not contain an <SQL routine body> that contains an SQL-dynamic statement.
- 26) Without Feature T653, “SQL-schema statements in external routines”, conforming SQL language shall not contain an <external routine name> that identifies a program in which an SQL-schema statement appears.
- 27) Without Feature T654, “SQL-dynamic statements in external routines”, conforming SQL language shall not contain an <external routine name> that identifies a program in which an SQL-dynamic statement appears.
- 28) Without Feature T655, “Cyclically dependent routines”, conforming SQL language shall not contain an <SQL routine body> that contains a <routine invocation> whose subject routine is generally dependent on the routine descriptor of the SQL-invoked routine specified by <SQL-invoked routine>.
- 29) Without Feature T272, “Enhanced savepoint management”, conforming SQL language shall not contain a <routine characteristics> that contains a <savepoint level indication>.
- 30) Without Feature B121, “Routine language Ada”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains ADA.

- 31) Without Feature B122, “Routine language C”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains C.
- 32) Without Feature B123, “Routine language COBOL”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains COBOL.
- 33) Without Feature B124, “Routine language Fortran”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains FORTRAN.
- 34) Without Feature B125, “Routine language MUMPS”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains M.
- 35) Without Feature B126, “Routine language Pascal”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains PASCAL.
- 36) Without Feature B127, “Routine language PL/I”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains PLI.
- 37) Without Feature B128, “Routine language SQL”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains SQL.

## 11.52 <alter routine statement>

This Subclause is modified by [Subclause 9.9](#), “<alter routine statement>”, in ISO/IEC 9075-13.

### Function

Alter a characteristic of an SQL-invoked routine.

### Format

```

<alter routine statement> ::=
 ALTER <specific routine designator>
 <alter routine characteristics> <alter routine behavior>

<alter routine characteristics> ::=
 <alter routine characteristic>...

<alter routine characteristic> ::=
 <language clause>
 | <parameter style clause>
 | <SQL-data access indication>
 | <null-call clause>
 | <returned result sets characteristic>
 | NAME <external routine name>

<alter routine behavior> ::=
 RESTRICT

```

### Syntax Rules

- 1) 13 Let *SR* be the SQL-invoked routine identified by the <specific routine designator> and let *SN* be the <specific name> of *SR*. The schema identified by the explicit or implicit <schema name> of *SN* shall include the descriptor of *SR*.
- 2) *SR* shall be a schema-level routine.
- 3) *SR* shall not be an SQL-invoked routine that is dependent on a user-defined type.  
NOTE 350 — “SQL-invoked routine dependent on a user-defined type” is defined in [Subclause 4.27](#), “SQL-invoked routines”.
- 4) If RESTRICT is specified, then:
  - a) *SR* shall not be the ordering function included in the descriptor of any user-defined type *UDT*.
  - b) *SR* shall not be the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any of the following:
    - i) The SQL routine body of any routine descriptor.
    - ii) The <query expression> of any view descriptor.
    - iii) The <search condition> of any constraint descriptor.
    - iv) The triggered action of any trigger descriptor.

- c) *SN* shall not be included in any of the following:
  - i) A group descriptor of any transform descriptor.
  - ii) A user-defined cast descriptor.
- 5) *SR* shall be an external routine.
- 6) *SR* shall not be an SQL-invoked method that is an overriding method and the set of overriding methods of *SR* shall be empty.
- 7) <alter routine characteristics> shall contain at most one <language clause>, at most one <parameter style clause>, at most one <SQL-data access indication>, at most one <null-call clause>, at most one <returned result sets characteristic>, and at most one <external routine name>.
- 8) If <returned result sets characteristic> is specified, then *SR* shall be an SQL-invoked procedure.
- 9) If <language clause> is specified, then:
  - a) <language clause> shall not specify SQL.
  - b) Depending on whether the <language clause> specifies ADA, C, COBOL, FORTRAN, M, PASCAL, or PLI, let the operative data type correspondences table be Table 16, “Data type correspondences for Ada”, Table 17, “Data type correspondences for C”, Table 18, “Data type correspondences for COBOL”, Table 19, “Data type correspondences for Fortran”, Table 20, “Data type correspondences for M”, Table 21, “Data type correspondences for Pascal”, or Table 22, “Data type correspondences for PL/I”, respectively. Refer to the two columns of the operative data type correspondences table as the “SQL data type” column and the “host data type column”.
  - c) Any <data type> in the effective SQL parameter list entry of *SR* shall specify a data type listed in the SQL data type column for which the corresponding row in the host data type column is not “None”.

NOTE 351 — “Effective SQL parameter list” is defined in Subclause 11.51, “<SQL-invoked routine>”.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of *SN*.

## General Rules

- 1) If *SR* is not a method, then the routine descriptor of *SR* is modified:
  - a) If <returned result sets characteristic> is specified, then the value of <maximum returned result sets>.
  - b) If <language clause> is specified, then the <language name> contained in the <language clause>.
  - c) If <external routine name> is specified, then the external name of the routine descriptor is <external routine name>.
  - d) If <parameter style clause> is specified, then the routine descriptor includes an indication of whether the parameter passing style is PARAMETER STYLE SQL or PARAMETER STYLE GENERAL.

- e) If the <SQL-data access indication> is specified, then an indication of whether the SQL-invoked routine's <SQL-data access indication> is READS SQL DATA, MODIFIES SQL DATA, CONTAINS SQL, or NO SQL.
  - f) If <null-call clause> is specified, then an indication of whether the SQL-invoked routine is a null-call function.
- 2) If *SR* is a method, then let *DMS* be the descriptor of the corresponding method specification. *DMS* is modified:
- a) If <language clause> is specified, then the <language name> contained in the <language clause>.
  - b) If <parameter style clause> is specified, then the method specification descriptor includes an indication of whether the parameter passing style is PARAMETER STYLE SQL or PARAMETER STYLE GENERAL.
  - c) If the <SQL-data access indication> is specified, then an indication of whether the SQL-invoked routine's <SQL-data access indication> is READS SQL DATA, MODIFIES SQL DATA, CONTAINS SQL, or NO SQL.
  - d) If <null-call clause> is specified, then an indication of whether the method should not be invoked if any argument is the null value.
- 3) If *SR* is a method, then the routine descriptor of *SR* is modified:
- a) If <external routine name> is specified, then the external name of the routine descriptor is <external routine name>. If <parameter style clause> is specified, then the method specification descriptor includes an indication of whether the parameter passing style is PARAMETER STYLE SQL or PARAMETER STYLE GENERAL.
- 4) The last-altered timestamp included in the routine descriptor of *SR* is the value of CURRENT\_TIMESTAMP.

## Conformance Rules

- 1) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <alter routine statement>.



## 11.53 <drop routine statement>

*This Subclause is modified by Subclause 9.25, “<drop routine statement>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 11.11, “<drop routine statement>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 9.10, “<drop routine statement>”, in ISO/IEC 9075-13.*

### Function

Destroy an SQL-invoked routine.

### Format

```
<drop routine statement> ::=
 DROP <specific routine designator> <drop behavior>
```

### Syntax Rules

- 1) Let *SR* be the SQL-invoked routine identified by the <specific routine designator> and let *SN* be the <specific name> of *SR*. The schema identified by the explicit or implicit <schema name> of *SN* shall include the descriptor of *SR*.
- 2) *SR* shall be a schema-level routine.
- 3) *SR* shall not be dependent on any user-defined type.  

NOTE 352 — The notion of an SQL-invoked routine being dependent on a user-defined type is defined in Subclause 4.27, “SQL-invoked routines”.
- 4) If RESTRICT is specified, then *SR* shall not be the ordering function included in the descriptor of any user-defined type *UDT*.
- 5) If RESTRICT is specified, then:
  - a) *SR* shall not be the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any of the following:
    - i) The SQL routine body of any routine descriptor.
    - ii) The <query expression> of any view descriptor.
    - iii) The <search condition> of any constraint descriptor.
    - iv) 04 The triggered action of any trigger descriptor.
  - b) *SN* shall not be included in any of the following:
    - i) A group descriptor of any transform descriptor.
    - ii) 09 A user-defined cast descriptor.

NOTE 353 — If CASCADE is specified, then such referencing objects will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.
- 6) Let the containing schema be the schema identified by the <schema name> explicitly or implicitly contained in *SN*.

## Access Rules

- 1) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of *SN*. The enabled authorization identifiers shall include *A*.

## General Rules

- 1) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE EXECUTE ON SPECIFIC ROUTINE SN FROM A CASCADE
```

- 2) 09 Let *DN* be the <user-defined type name> of a user-defined type whose descriptor includes *SN* in the group descriptor of any transform descriptor. Let *GN* be the <group name> of that group descriptor. The following <drop transform statement> is effectively executed without further Access Rule checking:

```
DROP TRANSFORM GN FOR DN CASCADE
```

- 3) Let *UDCD* be a user-defined cast descriptor that includes *SN* as its cast function. Let *SDT* be the source data type included in *UDCD*. Let *TDT* be the target data type included in *UDCD*. The following <drop user-defined cast statement> is effectively executed without further Access Rule checking:

```
DROP CAST (DN AS TD) CASCADE
```

- 4) If *SR* is the ordering function included in the descriptor of a user-defined type *UDT*, then let *UDTN* be a <path-resolved user-defined type name> that identifies *UDT*. The following <drop user-defined ordering statement> is effectively executed without further Access Rule checking:

```
DROP ORDERING FOR UDTN CASCADE
```

- 5) The descriptor of *SR* is destroyed.

## Conformance Rules

- 1) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop routine statement> that contains a <drop behavior> that contains CASCADE.
- 2) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <drop routine statement> that contains a <specific routine designator> that identifies a method.

## 11.54 <user-defined cast definition>

This Subclause is modified by Subclause 11.12, “<user-defined cast definition>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 12.8, “<user-defined cast definition>”, in ISO/IEC 9075-14.

### Function

Define a user-defined cast.

### Format


```
<user-defined cast definition> ::=
 CREATE CAST <left paren> <source data type> AS <target data type> <right paren>
 WITH <cast function>
 [AS ASSIGNMENT]

<cast function> ::=
 <specific routine designator>

<source data type> ::=
 <data type>

<target data type> ::=
 <data type>
```

### Syntax Rules

- 1)  Let *SDT* be the <source data type>. The data type identified by *SDT* is called the *source data type*.
- 2) Let *TDT* be the <target data type>. The data type identified by *TDT* is called the *target data type*.
- 3) There shall be no user-defined cast for *SDT* and *TDT*.
- 4) At least one of *SDT* or *TDT* shall contain a <schema-resolved user-defined type name> or a <reference type>.
- 5) If *SDT* contains a <schema-resolved user-defined type name>, then let *SSDT* be the schema that includes the descriptor of the user-defined type identified by *SDT*.
- 6) If *SDT* contains a <reference type>, then let *SSDT* be the schema that includes the descriptor of the referenced type of the reference type identified by *SDT*.
- 7) If *TDT* contains a <schema-resolved user-defined type name>, then let *STDT* be the schema that includes the descriptor of the user-defined type identified by *TDT*.
- 8) If *TDT* contains a <reference type>, then let *STDT* be the schema that includes the descriptor of the referenced type of the reference type identified by *TDT*.
- 9) If both *SDT* and *TDT* contain a <schema-resolved user-defined type name> or a <reference type>, then the <authorization identifier> that owns *SSDT* and the <authorization identifier> that owns *STDT* shall be equivalent.

**11.54 <user-defined cast definition>**

- 10) Let  $F$  be the SQL-invoked routine identified by <cast function>.  $F$  is called the *cast function* for source data type  $SDT$  and target data type  $TDT$ .
- a)  $F$  shall have exactly one SQL parameter, and its declared type shall be  $SDT$ .
  - b) The result data type of  $F$  shall be  $TDT$ .
  - c) The <authorization identifier> that owns  $SSDT$  or  $STDT$  (both, if both  $SDT$  and  $TDT$  are <schema-resolved user-defined type name>s) shall own the schema that includes the SQL-invoked routine descriptor of  $F$ .
  - d)  $F$  shall be deterministic.
  - e)  $F$  shall not possibly modify SQL-data.
  - f)  $F$  shall not possibly read SQL-data.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema that includes the routine descriptor of  $F$ .
- 2) If  $SDT$  contains a <schema-resolved user-defined type name> or a <reference type>, then the enabled authorization identifiers shall include the <authorization identifier> that owns  $SSDT$ .
- 3) If  $TDT$  contains a <schema-resolved user-defined type name> or a <reference type>, then the enabled authorization identifiers shall include the <authorization identifier> that owns  $STDT$ .

## General Rules

- 1) A user-defined cast descriptor  $CFD$  is created that describes the user-defined cast.  $CFD$  includes the name of the source data type, the name of the target data type, the specific name of the cast function, and, if and only if AS ASSIGNMENT is specified, an indication that the cast function is implicitly invocable.

## Conformance Rules

- 1) Without Feature S211, “User-defined cast functions”, conforming SQL language shall not contain a <user-defined cast definition>.

## 11.55 <drop user-defined cast statement>

This Subclause is modified by *Subclause 9.26*, “<drop user-defined cast statement>”, in ISO/IEC 9075-4.

### Function

Destroy a user-defined cast.

### Format

```
<drop user-defined cast statement> ::=
 DROP CAST <left paren> <source data type> AS <target data type> <right paren>
 <drop behavior>
```

### Syntax Rules

- 1) Let *SDT* be the <source data type> and let *TDT* be the <target data type>.
- 2) Let *CF* be the user-defined cast whose user-defined cast descriptor includes *SDT* as the source data type and *TDT* as the target data type.
- 3) Let *SN* be the specific name of the cast function *F* included in the user-defined cast descriptor of *CF*.
- 4) The schema identified by the <schema name> of *SN* shall include the descriptor of *F*.
- 5) Let *CS* be any <cast specification> such that:
  - a) The <value expression> of *CS* has declared type *P*.
  - b) The <cast target> of *CS* is either *TDT* or a domain with declared type *TDT*.
  - c) The type designator of *SDT* is in the type precedence of *P*.
  - d) No other data type *Q* whose type designator precedes *SDT* in the type precedence list of *P* such that there is a user-defined cast *CF<sub>q</sub>* whose user-defined cast descriptor includes *Q* as the source data type and *TDT* as the target data type.
- 6) Let *PS* be any SQL procedure statement that is dependent on *F*.
- 7) If RESTRICT is specified, then neither *CS* nor *PS* shall be generally contained in any of the following:
  - a) The SQL routine body of any routine descriptor.
  - b) The <query expression> of any view descriptor.
  - c) The <search condition> of any constraint descriptor.
  - d) 04 The triggered action of any trigger descriptor.

NOTE 354 — If CASCADE is specified, then such referencing objects will be dropped as specified in the General Rules of this Subclause.

## Access Rules

- 1) The enabled authorization identifier shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of *SN*.

## General Rules

- 1) Let *R* be any SQL-invoked routine that contains *CS* or *PS* in its SQL routine body. Let *SN* be the specific name of *R*. The following <drop routine statement> is effectively executed without further Access Rule checking:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

- 2) Let *V* be any view that contains *CS* or *PS* in its <query expression>. Let *VN* be the <table name> of *V*. The following <drop view statement> is effectively executed without further Access Rule checking:

```
DROP VIEW VN CASCADE
```

- 3) Let *T* be any table that contains *CS* or *PS* in the <search condition> of any constraint descriptor included in the table descriptor of *T*. Let *TN* be the <table name> of *T*. The following <drop table statement> is effectively executed without further Access Rule checking:

```
DROP TABLE TN CASCADE
```

- 4) Let *A* be any assertion that contains *CS* or *PS* in its <search condition>. Let *AN* be the <constraint name> of *A*. The following <drop assertion statement> is effectively executed without further Access Rule checking:

```
DROP ASSERTION AN CASCADE
```

- 5) Let *D* be any domain that contains *CS* or *PS* in the <search condition> of any constraint descriptor. Let *DN* be the <domain name> of *D*. The following <drop domain statement> is effectively executed without further Access Rule checking:

```
DROP DOMAIN DN CASCADE
```

- 6) Let *T* be any trigger whose trigger descriptor includes a triggered action that contains *CS* or *PS*. Let *TN* be the <trigger name> of *T*. The following <drop trigger statement> is effectively executed without further Access Rule checking:

```
DROP TRIGGER TN
```

- 7) The descriptor of *CF* is destroyed.

## Conformance Rules

- 1) Without Feature S211, “User-defined cast functions”, conforming SQL language shall not contain a <drop user-defined cast statement>.

## 11.56 <user-defined ordering definition>

*This Subclause is modified by Subclause 11.13, “<user-defined ordering definition>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 9.11, “<user-defined ordering definition>”, in ISO/IEC 9075-13.*

### Function

Define a user-defined ordering for a user-defined type.

### Format

```
<user-defined ordering definition> ::=
 CREATE ORDERING FOR <schema-resolved user-defined type name> <ordering form>

<ordering form> ::=
 <equals ordering form>
 | <full ordering form>

<equals ordering form> ::=
 EQUALS ONLY BY <ordering category>

<full ordering form> ::=
 ORDER FULL BY <ordering category>

13 <ordering category> ::=
 <relative category>
 | <map category>
 | <state category>

<relative category> ::=
 RELATIVE WITH <relative function specification>

<map category> ::=
 MAP WITH <map function specification>

<state category> ::=
 STATE [<specific name>]

<relative function specification> ::=
 <specific routine designator>

<map function specification> ::=
 <specific routine designator>
```

### Syntax Rules

- 1) Let *UDTN* be the <schema-resolved user-defined type name>. Let *UDT* be the user-defined type identified by *UDTN*.
- 2) The descriptor of *UDT* shall include an ordering form that specifies NONE.
- 3) If *UDT* is not a maximal supertype, then  
Case:

11.56 <user-defined ordering definition>

- a) If <equals ordering form> is specified, then the comparison form of every direct supertype of *UDT* shall be EQUALS.
- b) Otherwise, the comparison form of every direct supertype of *UDT* shall be FULL.
- 4) 13 If <relative category> or <state category> is specified, then *UDT* shall be a maximal supertype.
- 5) If <map category> is specified and *UDT* is not a maximal supertype, then the comparison category of every direct supertype of *UDT* shall be MAP.

NOTE 355 — The comparison categories of two user-defined types in the same subtype family shall be the same.

- 6) 13 Case:
  - a) If <state category> is specified, then
    - i) *UDT* shall not be a distinct type.
    - ii) EQUALS ONLY shall be specified.
    - iii) 09 The declared type of each attribute of *UDT* shall not be UDT-NC-ordered.
    - iv) Case:
      - 1) If <specific name> is specified, then let *SN* be <specific name>. If *SN* contains a <schema name>, then that <schema name> shall be equivalent to the <schema name> of *UDTN*.
      - 2) Otherwise, let *SN* be an implementation-dependent <specific name> whose <schema name> is equivalent to the <schema name> *S* of *UDTN*. This implementation-dependent <specific name> shall not be equivalent to the <specific name> of any other routine descriptor in the schema identified by *S*.
  - b) 13 Otherwise:
    - i) Let *F* be the SQL-invoked routine identified by the <specific routine designator> *SRD*.
    - ii) *F* shall be deterministic.
    - iii) *F* shall not possibly modify SQL-data.
- 7) If <relative function specification> is specified, then:
  - a) *F* shall have exactly two SQL parameters whose declared type is *UDT*.
  - b) *F* shall be an SQL-invoked regular function.
  - c) The result data type of *F* shall be INTEGER.

NOTE 356 — The Syntax Rules and General Rules of Subclause 8.2, “<comparison predicate>”, expect that the result of *F* is 0 (zero) if the two arguments are deemed to be equal, -1 if the first argument is deemed to be less than the second argument, and 1 (one) if the first argument is deemed to be greater than the second argument.
- 8) If <map function specification> is specified, then:
  - a) *F* shall have exactly one SQL parameter whose declared type is *UDT*.
  - b) The result data type of *F* shall be a predefined data type.
  - c) The result data type of *F* is an operand of an equality operation. The Syntax Rules of Subclause 9.9, “Equality operations”, apply.



- d) If FULL is specified, then the result data type of *F* is an operand of an ordering operation. The Syntax Rules of [Subclause 9.12](#), “Ordering operations”, apply.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema that includes the descriptor of *UDT* and the schema that includes the routine descriptor of *F*.

## General Rules

- 1) If <state category> is specified, then:
- Let  $C_1, \dots, C_n$  be the components of the representation of the user-defined type.
  - Let *SNUDT* be the <schema name> of the schema that includes the descriptor of *UDT*.
  - The following <SQL-invoked routine> is effectively executed:

```
CREATE FUNCTION SNUDT.EQUALS (UDT1 UDTN, UDT2 UDTN)
 RETURNS BOOLEAN
 SPECIFIC SN
 DETERMINISTIC
 CONTAINS SQL
 STATIC DISPATCH
 RETURN
 (TRUE AND
 UDT1.SPECIFICTYPE = UDT2.SPECIFICTYPE AND
 UDT1.C1 = UDT2.C1 AND
 ...
 UDT1.Cn = UDT2.Cn)
```

- 2) Case:
- If EQUALS is specified, then the ordering form in the user-defined type descriptor of *UDT* is set to EQUALS.
  - Otherwise, the ordering form in the user-defined type descriptor of *UDT* is set to FULL.
- 3) Case:
- If RELATIVE is specified, then the ordering category in the user-defined type descriptor of *UDT* is set to RELATIVE.
  - If MAP is specified, then the ordering category in the user-defined type descriptor of *UDT* is set to map.
  - 13 Otherwise, the ordering category in the user-defined type descriptor of *UDT* is set to STATE.
- 4) The <specific routine designator> identifying the ordering function, depending on the ordering category, in the descriptor of *UDT* is set to *SRD*.

## Conformance Rules

- 1) Without Feature S251, “User-defined orderings”, conforming SQL shall not contain a <user-defined ordering definition>.

NOTE 357 — If MAP is specified, then the Conformance Rules of Subclause 9.9, “Equality operations”, apply. If ORDER FULL BY MAP is specified, then the Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.

## 11.57 <drop user-defined ordering statement>

*This Subclause is modified by Subclause 9.20, “<drop user-defined ordering statement>”, in ISO/IEC 9075-4. This Subclause is modified by Subclause 9.12, “<drop user-defined ordering statement>”, in ISO/IEC 9075-13.*

### Function

Destroy a user-defined ordering method.

### Format

```
<drop user-defined ordering statement> ::=
 DROP ORDERING FOR <schema-resolved user-defined type name> <drop behavior>
```

### Syntax Rules

- 1) Let *UDTN* be the <schema-resolved user-defined type name>. Let *UDT* be the user-defined type identified by *UDTN*.
- 2) The descriptor of *UDT* shall include an ordering form that specifies EQUALS or FULL.
- 3) Let *OF* be the ordering function of *UDT*.
- 4) If RESTRICT is specified, then none of the following shall contain an operand of an equality operation, grouping operation or ordering operation whose declared type is some user-defined type *TI* whose comparison type is *UDT*:
  - a) The SQL routine body of any routine descriptor.
  - b) The <query expression> of any view descriptor.
  - c) The <search condition> of any constraint descriptor.
  - d) 04 The triggered action of any trigger descriptor.

NOTE 358 — If CASCADE is specified, then such referencing objects will be dropped as specified in the General Rules of this Subclause.

### Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of *UDTN*.

### General Rules

- 1) 04 Let *R* be any SQL-invoked routine whose SQL routine body contains an operand of an equality operation, grouping operation, or ordering operation whose declared type is some user-defined type *TI* whose comparison type is *UDT*. Let *SN* be the specific name of *R*. The following <drop routine statement> is effectively executed without further Access Rule checking:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

**11.57 <drop user-defined ordering statement>**

- 2) Let  $V$  be any view whose <query expression> contains an operand of an equality operation, grouping operation, or ordering operation whose declared type is some user-defined type  $TI$  whose comparison type is  $UDT$ . Let  $VN$  be the <table name> of  $V$ . The following <drop view statement> is effectively executed without further Access Rule checking:

```
DROP VIEW VN CASCADE
```

- 3) Let  $T$  be any table whose table descriptor includes a constraint descriptor of a constraint  $C$  whose <search condition> contains an operand of an equality operation, grouping operation, or ordering operation whose declared type is some user-defined type  $TI$  whose comparison type is  $UDT$ . Let  $TN$  be the <table name> of  $T$ . Let  $TCN$  be the <constraint name> of  $C$ . The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE TN DROP CONSTRAINT TCN CASCADE
```

- 4) Let  $A$  be any assertion whose <search condition> contains an operand of an equality operation, grouping operation, or ordering operation whose declared type is some user-defined type  $TI$  whose comparison type is  $UDT$ . Let  $AN$  be the <constraint name> of  $A$ . The following <drop assertion statement> is effectively executed without further Access Rule checking:

```
DROP ASSERTION AN CASCADE
```

- 5) Let  $D$  be any domain whose descriptor includes a constraint descriptor that includes an operand of an equality operation, grouping operation, or ordering operation whose declared type is some user-defined type  $TI$  whose comparison type is  $UDT$ . Let  $DN$  be the <domain name> of  $D$ . The following <drop domain statement> is effectively executed without further Access Rule checking:

```
DROP DOMAIN DN CASCADE
```

- 6) 04 Let  $T$  be any trigger whose triggered action contains an operand of an equality operation, grouping operation, or ordering operation whose declared type is some user-defined type  $TI$  whose comparison type is  $UDT$ . Let  $TN$  be the <trigger name> of  $T$ . The following <drop trigger statement> is effectively executed without further Access Rule checking:

```
DROP TRIGGER TN
```

- 7) In the descriptor of  $UDT$ , the ordering form is set to NONE and the ordering category is set to STATE. No ordering function is included in the descriptor of  $UDT$ .

## Conformance Rules

- 1) Without Feature S251, “User-defined orderings”, conforming SQL language shall not contain a <drop user-defined ordering statement>.

## 11.58 <transform definition>

### Function

Define one or more transform functions for a user-defined type.

### Format

```

<transform definition> ::=
 CREATE { TRANSFORM | TRANSFORMS } FOR
 <schema-resolved user-defined type name> <transform group>...

<transform group> ::=
 <group name> <left paren> <transform element list> <right paren>

<group name> ::=
 <identifier>

<transform element list> ::=
 <transform element> [<comma> <transform element>]

<transform element> ::=
 <to sql>
 | <from sql>

<to sql> ::=
 TO SQL WITH <to sql function>

<from sql> ::=
 FROM SQL WITH <from sql function>

<to sql function> ::=
 <specific routine designator>

<from sql function> ::=
 <specific routine designator>

```

### Syntax Rules

- 1) Let *TD* be the <transform definition>. Let *DTN* be the <schema-resolved user-defined type name> immediately contained in *TD*. Let *DT* be the data type identified by *DTN*. Let *SDT* be the schema that includes the descriptor of *DT*. Let *TRD* be the transform descriptor included in the data type descriptor of *DT*.
- 2) No two <transform group>s immediately contained in *TD* shall have the same <group name>.
- 3) The SQL-invoked function identified by <to sql function> is called the *to-sql function*. The SQL-invoked function identified by <from sql function> is called the *from-sql function*.
- 4) Let *n* be the number of <transform group>s immediately contained in *TD*. For *i* ranging from 1 to *n*:
  - a) Let *TG<sub>i</sub>* be the *i*-th <transform group> immediately contained in *TD*. Let *GN<sub>i</sub>* be the <group name> contained in *TG<sub>i</sub>*.

- b) Each of <to sql> and <from sql> immediately contained in  $TG_i$  shall be contained at most once in a <transform element list>.
- c) The SQL-invoked routines identified by <to sql function> and <from sql function> shall be SQL-invoked functions that are deterministic and do not possibly modify SQL-data.
- d)  $TRD$  shall not include a transform group descriptor  $GD$  that includes a group name that is equivalent to  $GN_i$ .
- e) Let  $SDTT$  be the set that includes every data type  $DTT_j$  that is either a proper supertype or a proper subtype of  $DT$  such that the transform descriptor included in the data type descriptor of  $DTT_j$  includes a group descriptor  $GDT_{j,k}$  that includes a group name that is equivalent to  $GN_i$ .  $SDTT$  shall be empty.
- f) If <to sql> is specified, then let  $TSF_i$  be the SQL-invoked function identified by <to sql function>.
  - i) Case:
    - 1) If  $TSF_i$  is an SQL-invoked method, then  $TSF_i$  shall have exactly two SQL parameters such that the declared type of the first SQL parameter is  $DT$  and the declared type of the second SQL parameter is a predefined data type. The result data type of  $TSF_i$  shall be  $DT$ .
    - 2) Otherwise,  $TSF_i$  shall have exactly one SQL parameter whose declared type is a predefined data type. The result data type of  $TSF_i$  shall be  $DT$ .
  - ii) If  $DT$  is a structured type and  $TSF_i$  is an SQL-invoked method, then  $TSF_i$  shall be a type-preserving function.
- g) If <from sql> is specified, then let  $FSF_i$  be the SQL-invoked function identified by <from sql function>.  $FSF_i$  shall have exactly one SQL parameter whose declared type is  $DT$ . The result data type of  $FSF_i$  shall be a predefined data type.
- h) If <to sql> and <from sql> are both specified, then
  - Case:
    - i) If  $TSF_i$  is an SQL-invoked method, then the result data type of  $FSF_i$  and the data type of the second SQL parameter of  $TSF_i$  shall be compatible.
    - ii) Otherwise, the result data type of  $FSF_i$  and the data type of the first SQL parameter of  $TSF_i$  shall be compatible.

## Access Rules

- 1) For  $i$  ranging from 1 to  $n$ , the enabled authorization identifiers shall include the <authorization identifier> that owns  $SDT$  and the schema that includes the routine descriptors of  $TSF_i$ , if any, and  $FSF_i$ , if any.

## General Rules

- 1) A <group name> specifies the group name that identifies a transform group.
- 2) For every  $TG_i$ ,  $1 \text{ (one)} \leq i \leq n$ :

- a) A new group descriptor  $GD_i$  is created that includes the <group name> immediately contained in  $TG_i$ .  $GD_i$  is included in the list of transform group descriptors included in  $TRD$ .
- b) If <to sql> is specified, then the specific name of the to-sql function in  $GD_i$  is set to  $TSF_i$ .
- c) If <from sql> is specified, then the specific name of the from-sql function in  $GD_i$  is set to  $FSF_i$ .

## Conformance Rules

- 1) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <transform definition>.

## 11.59 <alter transform statement>

### Function

Change the definition of one or more transform groups.

### Format

```

<alter transform statement> ::=
 ALTER { TRANSFORM | TRANSFORMS }
 FOR <schema-resolved user-defined type name> <alter group>...

<alter group> ::=
 <group name> <left paren> <alter transform action list> <right paren>

<alter transform action list> ::=
 <alter transform action> [{ <comma> <alter transform action> }...]

<alter transform action> ::=
 <add transform element list>
 | <drop transform element list>

```

### Syntax Rules

- 1) Let *DN* be the <schema-resolved user-defined type name> and let *D* be the data type identified by *DN*. The schema identified by the explicit or implicit schema name of *DN* shall include the data type descriptor of *D*. Let *S* be that schema. Let *TD* be the transform descriptor included in the data type descriptor of *D*.
- 2) The scope of *DN* is the entire <alter transform statement> *AT*.
- 3) Let *n* be the number of <group name>s contained in *AT*. For *i* ranging from 1 to *n*:
  - a) Let *GN<sub>i</sub>* be the *i*-th <group name> contained in *AT*.
  - b) For each *GN<sub>i</sub>*, there shall be a transform group descriptor included in *TD* whose group name is equivalent to *GN<sub>i</sub>*. Let *GD<sub>i</sub>* be this transform group descriptor.

### Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns *S*.

### General Rules

- 1) For *i* ranging from 1 to *n*, *GD<sub>i</sub>* is modified as specified by <alter transform action list>.

### Conformance Rules

- 1) Without Feature S242, “Alter transform statement”, conforming SQL language shall not contain an <alter transform statement>.



## 11.60 <add transform element list>

### Function

Add a transform element (<to sql> and/or <from sql>) to an existing transform group.

### Format

```
<add transform element list> ::=
 ADD <left paren> <transform element list> <right paren>
```

### Syntax Rules

- 1) Let *AD* be the <add transform element list>.
- 2) Let *DN* be the <schema-resolved user-defined type name> immediately contained in the containing <alter transform statement>. Let *D* be the user-defined type identified by *DN*. Let *TD* be the transform descriptor included in the data type descriptor of *D*.
- 3) Let *GD* be the transform group descriptor included in *TD* whose group name is equivalent to <group name> immediately contained in the containing <alter group>.
- 4) Each of <to sql> and <from sql> (immediately contained in *AD*) shall be contained at most once in the <transform element list>.
- 5) If *GD* includes a specific name of the to-sql function, then *AD* shall not contain <to sql>.
- 6) If *GD* includes a specific name of the from-sql function, then *AD* shall not contain <from sql>.
- 7) The SQL-invoked routine identified by either <to sql function> or <from sql function> shall be an SQL-invoked function that is deterministic and does not possibly modify SQL-data.
- 8) If <to sql> is specified, then let *TSF* be the SQL-invoked function identified by <to sql function>.
  - a) Case:
    - i) If *TSF* is an SQL-invoked method, then *TSF* shall have exactly two SQL parameters such that the declared type of the first SQL parameter is *D* and the declared type of the second SQL parameter is a predefined data type. The result data type of *TSF* shall be *D*.
    - ii) Otherwise, *TSF* shall have exactly one SQL parameter whose declared type is a predefined data type. The result data type of *TSF* shall be *D*.
  - b) If *D* is a structured type, then *TSF* shall be a type-preserving function.
  - c) If *GD* includes the specific name of a from-sql function, then let *FS* be the SQL-invoked function that is identified by this specific name.

Case:

    - i) If *TSF* is an SQL-invoked method, then the result data type of *FS* and the data type of the second SQL parameter of *TSF* shall be compatible.

- ii) Otherwise, the result data type of *FS* and the data type of the first SQL parameter of *TSF* shall be compatible.
- 9) If <from sql> is specified, then let *FSF* be the SQL-invoked function identified by <from sql function>.
- a) *FSF* shall have exactly one SQL parameter whose declared type is *D*. The result data type of *FSF* shall be a predefined data type.
  - b) If *GD* includes the specific name of a to-sql function, then let *TS* be the SQL-invoked routine that is identified by this specific name.

Case:

- i) If *TS* is an SQL-invoked method, then the result data type of *FSF* and the data type of the second SQL parameter of *TS* shall be compatible.
- ii) Otherwise, the result data type of *FSF* and the data type of the first SQL parameter of *TS* shall be compatible.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema that includes the routine descriptors of *TSF*, if any, and *FSF*, if any.

## General Rules

- 1) If <to sql> is specified, then the specific name of the to-sql function in *GD* is set to *TSF*.
- 2) If <from sql> is specified, then the specific name of the from-sql function in *GD* is set to *FSF*.

## Conformance Rules

*None.*

## 11.61 <drop transform element list>

### Function

Remove a transform element (<to sql> and/or <from sql>) from a transform group.

### Format

```
<drop transform element list> ::=
 DROP <left paren> <transform kind>
 [<comma> <transform kind>] <drop behavior> <right paren>

<transform kind> ::=
 TO SQL
 | FROM SQL
```

### Syntax Rules

- 1) Let *DN* be the <schema-resolved user-defined type name> immediately contained in the containing <alter transform statement>. Let *D* be the user-defined type identified by *DN*. Let *TD* be the transform descriptor included in the data type descriptor of *D*.
- 2) Let *GD* be the transform group descriptor included in *TD* whose group name is equivalent to <group name> immediately contained in the containing <alter group>.
- 3) Each of TO SQL and FROM SQL shall only be specified at most once in the <drop transform element list>.
- 4) If TO SQL is specified then *GD* shall include the specific name of a to-sql function. Let *TSF* be this function.
- 5) If FROM SQL is specified then *GD* shall include the specific name of a from-sql function. Let *FSF* be this function.
- 6) If RESTRICT is specified, then:
  - a) If TO SQL is specified, then there shall be no external routine that has an SQL parameter whose associated to-sql function is *TSF* nor shall there be an external function that has *TSF* as the to-sql function associated with the result.
  - b) If FROM SQL is specified, then there shall be no external routine that has an SQL parameter whose associated from-sql function is *FSF*.

### Access Rules

*None.*

### General Rules

- 1) If FROM SQL is specified, then:

11.61 <drop transform element list>

- a) Let *FSN* be the <specific name> of any external routine that has an SQL parameter whose associated from-sql function is *FSF*. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE *FSN* CASCADE

- b) The specific name of the from-sql function is removed from *GD*.

2) If TO SQL is specified, then:

- a) Let *TSN* be the <specific name> of any external routine that has an SQL parameter whose associated to-sql function is *TSF*. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE *TSN* CASCADE

- b) Let *RSN* be the <specific name> of any external function that has *TSF* as the to-sql function associated with the result. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE *RSN* CASCADE

- c) The specific name of the to-sql function is removed from *GD*.

## Conformance Rules

*None.*

## 11.62 <drop transform statement>

### Function

Remove one or more transform functions associated with a transform.

### Format

```
<drop transform statement> ::=
 DROP { TRANSFORM | TRANSFORMS } <transforms to be dropped>
 FOR <schema-resolved user-defined type name> <drop behavior>

<transforms to be dropped> ::=
 ALL
 | <transform group element>

<transform group element> ::=
 <group name>
```

### Syntax Rules

- 1) Let *DT* be the data type identified by <schema-resolved user-defined type name>. Let *SDT* be the schema that includes the descriptor of *DT*. Let *TRD* be the transform descriptor included in the data type descriptor of *DT*. Let *n* be the number of transform group descriptors in *TRD*.
- 2) If <transform group element> is specified, then *TRD* shall include a transform group descriptor *GD* that includes a group name that is equivalent to the <group name> immediately contained in <transform group element>.
- 3) If RESTRICT is specified, then  
Case:
  - a) If ALL is specified, then for *i* ranging from 1 (one) to *n*:
    - i) Let *GD<sub>i</sub>* be the *i*-th transform group descriptor included in *TRD*.
    - ii) If *GD<sub>i</sub>* includes the specific name of a from-sql function *FSF<sub>i</sub>* then there shall be no external routine that has an SQL parameter whose associated from-sql function is *FSF<sub>i</sub>*.
    - iii) If *GD<sub>i</sub>* includes the specific name of a to-sql function *TSF<sub>i</sub>* then there shall be no external routine that has an SQL parameter whose associated to-sql function is *TSF<sub>i</sub>* nor shall there be an external function that has *TSF<sub>i</sub>* as the to-sql function associated with the result.
  - b) Otherwise:
    - i) If *GD* includes the specific name of a from-sql function *FSF* then there shall be no external routine that has an SQL parameter whose associated from-sql function is *FSF*.
    - ii) If *GD* includes the specific name of a to-sql function *TSF* then there shall be no external routine that has an SQL parameter whose associated to-sql function is *TSF* nor shall there be an external function that has *TSF* as the to-sql function associated with the result.

## Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns *SDT*.

## General Rules

- 1) Case:
  - a) If ALL is specified, then, for  $i$  ranging from 1 (one) to  $n$ :
    - i) Let  $GD_i$  be the  $i$ -th transform group descriptor included in *TRD*.
    - ii) If  $GD_i$  includes the specific name of a from-sql function  $FSF_i$ , then let  $FSN$  be the <specific name> of any external routine that has an SQL parameter whose associated from-sql function is  $FSF_i$ . The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE  $FSN$  CASCADE

- iii) If  $GD_i$  includes the specific name of a to-sql function  $TSF_i$ , then:
      - 1) Let  $TSN$  be the <specific name> of any external routine that has an SQL parameter whose associated to-sql function is  $TSF_i$ . The following <drop routine statement> is effectively executed without further Access Rule checking:
    - 2) Let  $RSN$  be the <specific name> of any external function that has  $TSF_i$  as the to-sql function associated with the result. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE  $TSN$  CASCADE

DROP SPECIFIC ROUTINE  $RSN$  CASCADE

- iv)  $GD_i$  is removed from *TRD*.
- b) Otherwise:
  - i) If  $GD$  includes the specific name of a from-sql function  $FSF$ , then let  $FSN$  be the <specific name> of any external routine that has an SQL parameter whose associated from-sql function is  $FSF$ . The following <drop routine statement> is effectively executed without further Access Rule checking:
  - ii) If  $GD$  includes the specific name of a to-sql function  $TSF$ , then:
    - 1) Let  $TSN$  be the <specific name> of any external routine that has an SQL parameter whose associated to-sql function is  $TSF$ . The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE  $FSN$  CASCADE

DROP SPECIFIC ROUTINE  $TSN$  CASCADE

- 2) Let *RSN* be the <specific name> of any external function that has *TSF* as the to-sql function associated with the result. The following <drop routine statement> is effectively executed without further Access Rule checking:

DROP SPECIFIC ROUTINE *RSN* CASCADE

- iii) *GD* is removed from *TRD*.

## Conformance Rules

- 1) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <drop transform statement>.

## 11.63 <sequence generator definition>

### Function

Define an external sequence generator.

### Format

```

<sequence generator definition> ::=
 CREATE SEQUENCE <sequence generator name> [<sequence generator options>]

<sequence generator options> ::=
 <sequence generator option> ...

<sequence generator option> ::=
 <sequence generator data type option>
 | <common sequence generator options>

<common sequence generator options> ::=
 <common sequence generator option> ...

<common sequence generator option> ::=
 <sequence generator start with option>
 | <basic sequence generator option>

<basic sequence generator option> ::=
 <sequence generator increment by option>
 | <sequence generator maxvalue option>
 | <sequence generator minvalue option>
 | <sequence generator cycle option>

<sequence generator data type option> ::=
 AS <data type>

<sequence generator start with option> ::=
 START WITH <sequence generator start value>

<sequence generator start value> ::=
 <signed numeric literal>

<sequence generator increment by option> ::=
 INCREMENT BY <sequence generator increment>

<sequence generator increment> ::=
 <signed numeric literal>

<sequence generator maxvalue option> ::=
 MAXVALUE <sequence generator max value>
 | NO MAXVALUE

<sequence generator max value> ::=
 <signed numeric literal>

<sequence generator minvalue option> ::=
 MINVALUE <sequence generator min value>
 | NO MINVALUE

```



```

<sequence generator min value> ::=
 <signed numeric literal>

<sequence generator cycle option> ::=
 CYCLE
 | NO CYCLE

```

## Syntax Rules

- 1) Let *SEQ* be the sequence generator defined by the <sequence generator definition> *SEQD*.
- 2) If *SEQD* is contained in a <schema definition> *SD* and the <sequence generator name> *SQN* contains a <schema name>, then that <schema name> shall be equivalent to the implicit or explicit <schema name> of *SD*.
- 3) The schema identified by the explicit or implicit schema name of *SQN* shall not include a sequence generator descriptor whose sequence generator name is equivalent to *SQN*.
- 4) If *SEQD* is contained in a <schema definition>, then let *A* be the explicit or implicit <authorization identifier> of the <schema definition>. Otherwise, let *A* be the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of *SQN*.
- 5) Each of <sequence generator data type option>, <sequence generator start with option>, <sequence generator increment by option>, <sequence generator maxvalue option>, <sequence generator minvalue option>, and <sequence generator cycle option> shall be specified at most once.
- 6) If <sequence generator data type option> is specified, then <data type> shall be an exact numeric type *DT* with scale 0 (zero); otherwise, let *DT* be an implementation-defined exact numeric type with scale 0 (zero).
- 7) The Syntax Rules of [Subclause 9.24, “Creation of a sequence generator”](#), are applied with <common sequence generator options> as *OPTIONS* and *DT* as *DATA TYPE*.

## Access Rules

- 1) If a <sequence generator definition> is contained in an <SQL-client module definition>, then the enabled authorization identifiers shall include *A*.

## General Rules

- 1) The General Rules of [Subclause 9.24, “Creation of a sequence generator”](#), are applied with <common sequence generator options> as *OPTIONS* and *DT* as *DATA TYPE*, yielding a sequence generator descriptor *SEQDS*. The sequence generator name included in *SEQDS* is set to *SQN*.
- 2) A privilege descriptor is created that defines the USAGE privilege on *SEQ* to *A*. This privilege is grantable. The grantor for this privilege descriptor is set to the special grantor value “\_SYSTEM”.

## Conformance Rules

- 1) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <sequence generator definition>.

## 11.64 <alter sequence generator statement>

### Function

Change the definition of an external sequence generator.

### Format

```

<alter sequence generator statement> ::=
 ALTER SEQUENCE <sequence generator name> <alter sequence generator options>

<alter sequence generator options> ::=
 <alter sequence generator option>...

<alter sequence generator option> ::=
 <alter sequence generator restart option>
 | <basic sequence generator option>

<alter sequence generator restart option> ::=
 RESTART WITH <sequence generator restart value>

<sequence generator restart value> ::=
 <signed numeric literal>

```

### Syntax Rules

- 1) Let *SEQ* be the sequence generator descriptor identified by the <sequence generator name> *SQN*. Let *DT* be the data type of *SEQ*.
- 2) The schema identified by the explicit or implicit schema name of *SQN* shall include *SEQ*.
- 3) The scope of *SQN* is the <alter sequence generator statement>.
- 4) The Syntax Rules of [Subclause 9.25, “Altering a sequence generator”](#), are applied with <alter sequence generator options> as *OPTIONS* and *SEQ* as *SEQUENCE*.

### Access Rules

- 1) The enabled authorization identifiers shall include the <authorization identifier> that owns the schema identified by the explicit or implicit schema name of *SQN*.

### General Rules

- 1) The General Rules of [Subclause 9.25, “Altering a sequence generator”](#), are applied with <alter sequence generator options> as *OPTIONS* and *SEQ* as *SEQUENCE*.

### Conformance Rules

- 1) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain an <alter sequence generator statement>.

## 11.65 <drop sequence generator statement>

### Function

Destroy an external sequence generator.

### Format

```
<drop sequence generator statement> ::=
 DROP SEQUENCE <sequence generator name> <drop behavior>
```

### Syntax Rules

- 1) Let *SEQ* be the sequence generator identified by the <sequence generator name> *SQN*.
- 2) The schema identified by the explicit or implicit schema name of *SQN* shall include the descriptor of *SEQ*.
- 3) If RESTRICT is specified, then *SEQ* shall not be referenced in any of the following:
  - a) The SQL routine body of any routine descriptor.
  - b) The triggered action of any trigger descriptor.

NOTE 359 — If CASCADE is specified, then such referenced objects will be dropped by the execution of the <revoke statement> specified in the General Rules of this Subclause.
- 4) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of the sequence generator identified by *SQN*.

### Access Rules

- 1) The enabled authorization identifiers shall include *A*.

### General Rules

- 1) The following <revoke statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```
REVOKE USAGE ON SEQUENCE SQN FROM A CASCADE
```

- 2) The descriptor of *SEQ* is destroyed.

### Conformance Rules

- 1) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <drop sequence generator statement>.

*(Blank page)*

## 12 Access control

### 12.1 <grant statement>

*This Subclause is modified by Subclause 10.1, “<grant statement>”, in ISO/IEC 9075-4.*

#### Function

Define privileges and role authorizations.

#### Format

```
<grant statement> ::=
 <grant privilege statement>
 | <grant role statement>
```

#### Syntax Rules

*None.*

#### Access Rules

*None.*

#### General Rules

- 1) For every involved grantee *G* and for every domain *DI* owned by *G*, if all of the following are true:
  - a) The applicable privileges for *G* include the grantable REFERENCES privilege on every column referenced in the <search condition> *SC* included in a domain constraint descriptor included in the domain descriptor of *DI*.
  - b) The applicable privileges for *G* include the grantable EXECUTE privileges on all SQL-invoked routines that are subject routines of <routine invocation>s contained in *SC*.
  - c) The applicable privileges for *G* include the grantable SELECT privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in *SC* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the method identified by the <method name> of *MR* included in a domain constraint descriptor included in the domain descriptor of *DI*.
  - d) The applicable privileges for *G* include the grantable SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of every <reference resolution> contained in *SC*.

## 12.1 &lt;grant statement&gt;

- e) The applicable privileges for *G* include the grantable USAGE privilege on all domains, character sets, collations, and transliterations whose <domain name>s, <character set name>s, <collation name>s, and <transliteration name>s, respectively, are included in the domain descriptor of *D1*.

then for every privilege descriptor with <action> USAGE, a grantor of “\_SYSTEM”, object *D1*, and grantee *G* that is not grantable, the following <grant statement> is effectively executed with a current user identifier of “\_SYSTEM” and without further Access Rule checking:

```
GRANT USAGE ON DOMAIN D1
TO G WITH
GRANT OPTION
```

- 2) For every involved grantee *G* and for every collation *C1* owned by *G*, if the applicable privileges for *G* include a grantable USAGE privilege for the character set name included in the collation descriptor of *C1* and a grantable USAGE privilege for the transliteration name, if any, included in the collation descriptor of *C1*, then for every privilege descriptor with <action> USAGE, a grantor of “\_SYSTEM”, object of *C1*, and grantee *G* that is not grantable, the following <grant statement> is effectively executed with a current user identifier of “\_SYSTEM” and without further Access Rule checking:

```
GRANT USAGE ON
 COLLATION C1
TO G
WITH GRANT OPTION
```

- 3) For every involved grantee *G* and for every transliteration *T1* owned by *G*, if the applicable privileges for *G* contain a grantable USAGE privilege for every character set identified by a <character set specification> contained in the <transliteration definition> of *T1*, then for every privilege descriptor with <action> *P*, a grantor of “\_SYSTEM”, object of *T1*, and grantee *G* that is not grantable, the following <grant statement> is effectively executed as though the current user identifier were “\_SYSTEM” and without further Access Rule checking:

```
GRANT P
 ON TRANSLATION T1
TO G
WITH GRANT OPTION
```

- 4) For every table *T* specified by some involved privilege descriptor and for each view *V* owned by some involved grantee *G* such that *T* or some column *CT* of *T* is referenced in the <query expression> *QE* of *V*, or *T* is a supertable of the scoped table of a <reference resolution> contained in *QE*, let *RT<sub>i</sub>*, for *i* ranging from 1 (one) to the number of tables identified by the <table reference>s contained in *QE*, be the <table name>s of those tables. For every column *CV* of *V*:
- Let *CRT<sub>i,j</sub>*, for *j* ranging from 1 (one) to the number of columns of *RT<sub>i</sub>* that are underlying columns of *CV*, be the <column name>s of those columns.
  - If, following successful execution of the <grant statement>, all of the following are true:
    - The applicable privileges for *G* include grantable SELECT privileges on all of the columns *CRT<sub>i,j</sub>*.
    - The applicable privileges for *G* include grantable EXECUTE privileges on all SQL-invoked routines that are subject routines of <routine invocation>s contained in *QE*.

- iii) The applicable privileges for  $G$  include grantable SELECT privilege on every table  $T1$  and every method  $M$  such that there is a <method reference>.  $MR$  contained in  $QE$  such that  $T1$  is in the scope of the <value expression primary> of  $MR$  and  $M$  is the method identified by the <method name> of  $MR$ .
- iv) The applicable privileges for  $G$  include grantable SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of every <reference resolution> that is contained in  $QE$ .

then the following <grant statement> is effectively executed as though the current user identifier were “\_SYSTEM” and without further Access Rule checking:

```
GRANT SELECT (CV)
 ON V
 TO G
 WITH GRANT OPTION
```

- c) If, following successful execution of the <grant statement>, the applicable privileges for  $G$  will include REFERENCES( $CRT_{i,j}$ ) for all  $i$  and for all  $j$ , and will include a REFERENCES privilege on some column of  $RT_i$  for all  $i$ , then:

- i) Case:

- 1) If all of the following are true, then let  $WGO$  be “WITH GRANT OPTION”.

- A) The applicable privileges for  $G$  will include grantable REFERENCES( $CRT_{i,j}$ ) for all  $i$  and for all  $j$ , and will include a grantable REFERENCES privilege on some column of  $RT_i$  for all  $i$ .
- B) The applicable privileges for  $G$  include grantable EXECUTE privileges on all SQL-invoked routines that are subject routines of <routine invocation>s contained in  $QE$ .
- C) The applicable privileges for  $G$  include grantable SELECT privilege on every table  $T1$  and every method  $M$  such that there is a <method reference>.  $MR$  contained in  $QE$  such that  $T1$  is in the scope of the <value expression primary> of  $MR$  and  $M$  is the method identified by the <method name> of  $MR$ .
- D) The applicable privileges for  $G$  include grantable SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of every <reference resolution> that is contained in  $QE$ .

- 2) Otherwise, let  $WGO$  be a zero-length string.

- ii) The following <grant statement> is effectively executed as though the current user identifier were “\_SYSTEM” and without further Access Rule checking:

```
GRANT REFERENCES (CV)
 ON V
 TO G
 WGO
```

- d) If, following successful execution of the <grant statement>, the applicable privileges for  $G$  include grantable SELECT privilege on every column of  $V$ , then the following <grant statement> is effectively executed as though the current user identifier were “\_SYSTEM” and without further Access Rule checking:

## 12.1 &lt;grant statement&gt;

```

GRANT SELECT
 ON V
 TO G
 WITH GRANT OPTION

```

- e) Following successful execution of the <grant statement>,

Case:

- i) If the applicable privileges for *G* include grantable REFERENCES privilege on every column of *V*, then let *WGO* be “WITH GRANT OPTION”.
- ii) If the applicable privileges for *G* include REFERENCES privilege on every column of *V*, then let *WGO* be a zero-length string.

The following <grant statement> is effectively executed as though the current user identifier were “\_SYSTEM” and without further Access Rule checking:

```

GRANT REFERENCES
 ON V
 TO G
 WGO

```

- 5) Following the successful execution of the <grant statement>, for every table *T* specified by some involved privilege descriptor and for every effectively updatable view *V* owned by some grantee *G* such that *T* is some leaf underlying table of the <query expression> of *V*, the General Rules of [Subclause 10.11](#), “[Determination of view privileges](#)”, are applied with *V* as the *VIEW*.
- 6) For every involved grantee *G* and for every referenceable view *V*, named *VN*, owned by *G*, if following the successful execution of the <grant statement>, the applicable privileges for *G* include grantable UNDER privilege on the direct supertable of *V*, then the following <grant statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```

GRANT UNDER
 ON VN
 TO G
 WITH GRANT OPTION

```

- 7) For every involved grantee *G* and for every schema-level SQL-invoked routine *RI* owned by *G*, if the applicable privileges for *G* contain all of the privileges necessary to successfully execute every <SQL procedure statement> contained in the <routine body> of *RI* are grantable, then for every privilege descriptor with <action> EXECUTE, a grantor of “\_SYSTEM”, object of *RI*, and grantee *G* that is not grantable, the following <grant statement> is effectively executed with a current authorization identifier of “\_SYSTEM” and without further Access Rule checking:

```

GRANT EXECUTE
 ON RI
 TO G
 WITH GRANT OPTION

```

NOTE 360 — The privileges necessary include the EXECUTE privilege on every subject routine of every <routine invocation> contained in the <SQL procedure statement>.



- 8) If two privilege descriptors are identical except that one indicates that the privilege is grantable and the other indicates that the privilege is not grantable, then both privilege descriptors are set to indicate that the privilege is grantable.
- 9) If two privilege descriptors are identical except that one indicates WITH HIERARCHY OPTION and the other does not, then both privilege descriptors are set to indicate that the privilege has the WITH HIERARCHY OPTION.
- 10) Redundant duplicate privilege descriptors are removed from the collection of all privilege descriptors.
- 11) Redundant duplicate view privilege dependency descriptors are removed from the collection of all view privilege dependency descriptors.

## Conformance Rules

*None.*

## 12.2 <grant privilege statement>

This Subclause is modified by *Subclause 10.1*, “<grant privilege statement>”, in ISO/IEC 9075-13.

### Function

Define privileges.

### Format

```
<grant privilege statement> ::=
 GRANT <privileges> TO <grantee> [{ <comma> <grantee> }...]
 [WITH HIERARCHY OPTION]
 [WITH GRANT OPTION]
 [GRANTED BY <grantor>]
```

### Syntax Rules

- 1) Let *O* be the object identified by the <object name> contained in <privileges>.
- 2) The schema identified by the explicit or implicit qualifier of the <object name> shall include the descriptor of *O*.
- 3) If WITH HIERARCHY OPTION is specified, then:
  - a) <privileges> shall specify an <action> of SELECT without a <privilege column list> and without a <privilege method list>.
  - b) *O* shall be a table of a structured type.

### Access Rules

*None.*

### General Rules

- 1) Case:
  - a) If GRANTED BY is omitted, then let *G* be OMITTED.
  - b) Otherwise, let *G* be <grantor>.
- 2) Let *A* be the result of applying the General Rules of *Subclause 12.8*, “Grantor determination”, with *G* as GRANTOR.
- 3) If the applicable privileges for *A* do not include a privilege identifying *O*, then an exception condition is raised: *privilege not granted*.
- 4) A set of privilege descriptors is identified. The privilege descriptors identified are those defining, for each <action> explicitly or implicitly in <privileges>, that <action> on *O* held by *A* with grant option.

- 5) The <object privileges> specify one or more privileges on the object identified by the <object name>.
- 6) For every identified privilege descriptor *IPD*, a privilege descriptor is created for each <grantee>, that specifies grantee <grantee>, action <action>, object *O*, and grantor *A*. Let *CPD* be the set of privilege descriptors created.
- 7) For every privilege descriptor in *CPD* whose action is INSERT, UPDATE, or REFERENCES without a column name, privilege descriptors are also created and added to *CPD* for each column *C* in *O* for which *A* holds the corresponding privilege with grant option. For each such column, a privilege descriptor is created that specifies the identical <grantee>, the identical <action>, object *C*, and grantor *A*.
- 8) For every privilege descriptor in *CPD* whose action is SELECT without a column name or method name, privilege descriptors are also created and added to *CPD* for each column *C* in *O* for which *A* holds the corresponding privilege with grant option. For each such column, a privilege descriptor is created that specifies the identical <grantee>, the identical <action>, object *C*, and grantor *A*.
- 9) For every privilege descriptor in *CPD* whose action is SELECT without a column name or method name, if the table *T* identified by the object of the privilege descriptor is a table of a structured type *TY*, then table/method privilege descriptors are also created and added to *CPD* for each method *M* of *TY* for which *A* holds the corresponding privilege with grant option. For each such method, a table/method privilege descriptor is created that specifies the identical <grantee>, the identical <action>, object consisting of the pair of table *T* and method *M*, and grantor *A*.
- 10) If WITH GRANT OPTION was specified, then each privilege descriptor also indicates that the privilege is grantable.
- 11) Let *SWH* be the set of privilege descriptors in *CPD* whose action is SELECT WITH HIERARCHY OPTION. Let *ST* be the set of subtables of *O*. For every table *T* in *ST* and for every privilege descriptor in *SWH* grantee *G*, and grantor *A*,  
Case:
  - a) If the privilege is grantable, then let *WGO* be “WITH GRANT OPTION”.
  - b) Otherwise, let *WGO* be the zero-length string.

The following <grant statement> is effectively executed without further Access Rule checking:

```
GRANT SELECT
 ON T
 TO G
 WGO
 GRANTED BY A
```

- 12) For every combination of <grantee> and <action> on *O* specified in <privileges>, if there is no corresponding privilege descriptor in *CPD*, then a completion condition is raised: *warning — privilege not granted*.
- 13) If ALL PRIVILEGES was specified, then for each grantee *G*, if there is no privilege descriptor in *CPD* specifying grantee *G*, then a completion condition is raised: *warning — privilege not granted*.
- 14) The *set of involved privilege descriptors* is defined to be *CPD*.
- 15) The *set of involved grantees* is defined as the set of specified <grantee>s.

## **Conformance Rules**

- 1) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <specific routine designator> contained in a <grant privilege statement> that identifies a method.
- 2) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <grant privilege statement> that contains WITH HIERARCHY OPTION.

## 12.3 <privileges>

*This Subclause is modified by Subclause 10.2, “<privileges>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 13.1, “<privileges>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 10.2, “<privileges>”, in ISO/IEC 9075-13.*

### Function

Specify privileges.

### Format

```

<privileges> ::=
 <object privileges> ON <object name>

04 09 13 <object name> ::=
 [TABLE] <table name>
 | DOMAIN <domain name>
 | COLLATION <collation name>
 | CHARACTER SET <character set name>
 | TRANSLATION <transliteration name>
 | TYPE <schema-resolved user-defined type name>
 | SEQUENCE <sequence generator name>
 | <specific routine designator>

<object privileges> ::=
 ALL PRIVILEGES
 | <action> [{ <comma> <action> }...]

<action> ::=
 SELECT
 | SELECT <left paren> <privilege column list> <right paren>
 | SELECT <left paren> <privilege method list> <right paren>
 | DELETE
 | INSERT [<left paren> <privilege column list> <right paren>]
 | UPDATE [<left paren> <privilege column list> <right paren>]
 | REFERENCES [<left paren> <privilege column list> <right paren>]
 | USAGE
 | TRIGGER
 | UNDER
 | EXECUTE

<privilege method list> ::=
 <specific routine designator> [{ <comma> <specific routine designator> }...]

<privilege column list> ::=
 <column name list>

<grantee> ::=
 PUBLIC
 | <authorization identifier>

<grantor> ::=
 CURRENT_USER
 | CURRENT_ROLE

```

## Syntax Rules

- 1) ALL PRIVILEGES is equivalent to the specification of all of the privileges on <object name> for which the <grantor> has grantable privilege descriptors.
- 2) If the <object name> of the <grant statement> or <revoke statement> specifying <privileges> specifies <table name>, then let *T* be the table identified by that <table name>. *T* shall not be a declared local temporary table.
- 3) 09 13 If <object name> specifies a <domain name>, <collation name>, <character set name>, <transliteration name>, <schema-resolved user-defined type name>, or <sequence generator name>, then <privileges> may specify USAGE. Otherwise, USAGE shall not be specified.
- 4) If <object name> specifies a <table name> that identifies a base table or a viewed table, then <privileges> may specify TRIGGER; otherwise, TRIGGER shall not be specified.
- 5) If <object name> specifies a <schema-resolved user-defined type name> that identifies a structured type or specifies a <table name>, then <privileges> may specify UNDER; otherwise, UNDER shall not be specified.
- 6) If *T* is a temporary table, then <privileges> shall specify ALL PRIVILEGES.
- 7) 04 If the object identified by <object name> of the <grant statement> or <revoke statement> is an SQL-invoked routine, then <privileges> may specify EXECUTE; otherwise, EXECUTE shall not be specified.
- 8) The <object privileges> specify one or more privileges on the object identified by <object name>.
- 9) Each <column name> in a <privilege column list> shall identify a column of *T*.
- 10) If <privilege method list> is specified, then <object name> shall specify a <table name> that identifies a table of a structured type *TY* and each <specific routine designator> in the <privilege method list> shall identify a method of *TY*.
- 11) UPDATE (<privilege column list>) is equivalent to the specification of UPDATE (<column name>) for each <column name> in <privilege column list>. INSERT (<privilege column list>) is equivalent to the specification of INSERT (<column name>) for each <column name> in <privilege column list>. REFERENCES (<privilege column list>) is equivalent to the specification of REFERENCES (<column name>) for each <column name> in <privilege column list>. SELECT (<privilege column list>) is equivalent to the specification of SELECT (<column name>) for each <column name> in <privilege column list>. SELECT (<privilege method list>) is equivalent to the specification of SELECT (<specific routine designator>) for each <specific routine designator> in <privilege method list>.

## Access Rules

*None.*

## General Rules

- 1) A <grantee> of PUBLIC denotes at all times a list of <grantee>s containing all of the <authorization identifier>s in the SQL-environment.
- 2) SELECT (<column name>) specifies the SELECT privilege on the indicated column and implies one or more column privilege descriptors.

- 3) SELECT (<specific routine designator>) specifies the SELECT privilege on the indicated method for the table identified by <object name> and implies one or more table/method privilege descriptors.
- 4) SELECT with neither <privilege column list> nor <privilege method list> specifies the SELECT privilege on all columns of *T* including any columns subsequently added to *T* and implies a table privilege descriptor and one or more column privilege descriptors. If *T* is a table of a structured type *TY*, then SELECT also specifies the SELECT privilege on all methods of the type *TY*, including any methods subsequently added to the type *TY*, and implies one or more table/method privilege descriptors.
- 5) UPDATE (<column name>) specifies the UPDATE privilege on the indicated column and implies one or more column privilege descriptors. If the <privilege column list> is omitted, then UPDATE specifies the UPDATE privilege on all columns of *T*, including any column subsequently added to *T* and implies a table privilege descriptor and one or more column privilege descriptors.
- 6) INSERT (<column name>) specifies the INSERT privilege on the indicated column and implies one or more column privilege descriptors. If the <privilege column list> is omitted, then INSERT specifies the INSERT privilege on all columns of *T*, including any column subsequently added to *T* and implies a table privilege descriptor and one or more column privilege descriptors.
- 7) REFERENCES (<column name>) specifies the REFERENCES privilege on the indicated column and implies one or more column privilege descriptors. If the <privilege column list> is omitted, then REFERENCES specifies the REFERENCES privilege on all columns of *T*, including any column subsequently added to *T* and implies a table privilege descriptor and one or more column privilege descriptors.
- 8) *B* has the WITH ADMIN OPTION on a role if a role authorization descriptor identifies the role as granted to *B* WITH ADMIN OPTION or a role authorization descriptor identifies it as granted WITH ADMIN OPTION to another applicable role for *B*.

## Conformance Rules

- 1) Without Feature T332, “Extended roles”, conforming SQL language shall not contain a <grantor>.
- 2) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain an <action> that contains TRIGGER.
- 3) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <privileges> that contains an <action> that contains UNDER and that contains an <object name> that contains a <table name>.
- 4) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <privileges> that contains an <action> that contains UNDER and that contains an <object name> that contains a <schema-resolved user-defined type name> that identifies a structured type.
- 5) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <privileges> that contains an <action> that contains USAGE and that contains an <object name> that contains a <schema-resolved user-defined type name> that identifies a structured type.
- 6) Without Feature T281, “SELECT privilege with column granularity”, in conforming SQL language, an <action> that contains SELECT shall not contain a <privilege column list>.
- 7) Without Feature F731, “INSERT column privileges”, in conforming SQL language, an <action> that contains INSERT shall not contain a <privilege column list>.
- 8) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <privilege method list>.

## 12.4 <role definition>

### Function

Define a role.

### Format

```
<role definition> ::=
 CREATE ROLE <role name> [WITH ADMIN <grantor>]
```

### Syntax Rules

- 1) The specified <role name> shall not be equivalent to any other <authorization identifier> in the SQL-environment.

### Access Rules

- 1) The privileges necessary to execute the <role definition> are implementation-defined.

### General Rules

- 1) A role descriptor whose role name is <role name> is created in the SQL-environment.
- 2) Case:
  - a) If WITH ADMIN is omitted, then let *G* be OMITTED.
  - b) Otherwise, let *G* be <grantor>.
- 3) Let *A* be the result of applying the General Rules of [Subclause 12.8, “Grantor determination”](#), with *G* as *GRANTOR*.
- 4) A grantable role authorization descriptor is created whose role name is <role name>, whose grantor of “\_SYSTEM”, and whose grantee is *A*.

### Conformance Rules

- 1) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <role definition>.
- 2) Without Feature T332, “Extended roles”, conforming SQL language shall not contain a <role definition> that immediately contains WITH ADMIN.



## 12.5 <grant role statement>

### Function

Define role authorizations.

### Format

```
<grant role statement> ::=
 GRANT <role granted> [{ <comma> <role granted> }...]
 TO <grantee> [{ <comma> <grantee> }...]
 [WITH ADMIN OPTION]
 [GRANTED BY <grantor>]

<role granted> ::=
 <role name>
```

### Syntax Rules

- 1) No role identified by a specified <grantee> shall be applicable for any role identified by a specified <role granted>.

NOTE 361 — That is, no cycles of role authorizations are allowed.

### Access Rules

*None.*

### General Rules

- 1) Case:
  - a) If WITH ADMIN is omitted, then let *G* be OMITTED.
  - b) Otherwise, let *G* be <grantor>.
- 2) Let *A* be the result of applying the General Rules of [Subclause 12.8, “Grantor determination”](#), with *G* as *GRANTOR*.
- 3) For each <role granted> *R*, if no grantable role authorization descriptor exists whose role name is *R* and whose grantee is *A* or an applicable role for *A*, then an exception condition is raised: *invalid role specification*.
- 4) For each <grantee> *GEE*, for each <role granted> *R*, a role authorization descriptor is created with role name *R*, grantee *GEE*, and grantor *A*.
- 5) If WITH ADMIN OPTION is specified, then each role authorization descriptor is grantable.
- 6) If two role authorization descriptors are identical except that one is grantable and the other is not, then both role authorization descriptors are set to indicate that the role authorization is grantable.
- 7) Redundant duplicate role authorization descriptors are destroyed.

**12.5 <grant role statement>**

- 8) The *set of involved privilege descriptors* is the union of the sets of privilege descriptors corresponding to the applicable privileges for every <role granted> specified.
- 9) The *set of involved grantees* is the union of the set of <grantee>s and the set of <role name>s for which at least one of the <role name>s that is possibly specified as a <grantee> is applicable.

**Conformance Rules**

- 1) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <grant role statement>.

## 12.6 <drop role statement>

### Function

Destroy a role.

### Format

```
<drop role statement> ::=
 DROP ROLE <role name>
```

### Syntax Rules

- 1) Let  $R$  be the role identified by the specified <role name>.

### Access Rules

- 1) There shall exist at least one grantable role authorization descriptor whose role name is  $R$  and whose grantee is an enabled authorization identifier.

### General Rules

- 1) Let  $A$  be any <authorization identifier> identified by a role authorization descriptor as having been granted to  $R$ .
- 2) The following <revoke role statement> is effectively executed without further Access Rule checking:

```
REVOKE R FROM A
```

- 3) The descriptor of  $R$  is destroyed.

### Conformance Rules

- 1) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <drop role statement>.

## 12.7 <revoke statement>

*This Subclause is modified by Subclause 10.3, “<revoke statement>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 13.2, “<revoke statement>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 10.3, “<revoke statement>”, in ISO/IEC 9075-13.*

### Function

Destroy privileges and role authorizations.

### Format

```

<revoke statement> ::=
 <revoke privilege statement>
 | <revoke role statement>

<revoke privilege statement> ::=
 REVOKE [<revoke option extension>] <privileges>
 FROM <grantee> [{ <comma> <grantee> }...]
 [GRANTED BY <grantor>]
 <drop behavior>

<revoke option extension> ::=
 GRANT OPTION FOR
 | HIERARCHY OPTION FOR

<revoke role statement> ::=
 REVOKE [ADMIN OPTION FOR] <role revoked> [{ <comma> <role revoked> }...]
 FROM <grantee> [{ <comma> <grantee> }...]
 [GRANTED BY <grantor>]
 <drop behavior>

<role revoked> ::=
 <role name>

```

### Syntax Rules

- 1) Let *O* be the object identified by the <object name> contained in <privileges>. If *O* is a table *T*, then let *S* be the set of subtables of *O*. If *T* is a table of a structured type, then let *TY* be that type.
- 2) If WITH HIERARCHY OPTION is specified, the <privileges> shall specify an <action> of SELECT without a <privilege column list> and without a <privilege method list> and *O* shall be a table of a structured type.
- 3) SELECT is equivalent to specifying both the SELECT table privilege and SELECT (<privilege column list>) for all columns of <table name>. If *T* is a table of a structured type *TY*, then SELECT also specifies SELECT (<privilege column list>) for all columns inherited from *T* in each of the subtables of *T*, and SELECT (<privilege method list>) for all methods of *TY* in each of the subtables of *T*.
- 4) INSERT is equivalent to specifying both the INSERT table privilege and INSERT (<privilege column list>) for all columns of <table name>.

- 5) UPDATE is equivalent to specifying both the UPDATE table privilege and UPDATE (<privilege column list>) for all columns of <table name>, as well as UPDATE (<privilege column list>) for all columns inherited from *T* in each of the subtables of *T*.
- 6) REFERENCES is equivalent to specifying both the REFERENCES table privilege and REFERENCES (<privilege column list>) for all columns of <table name>, as well as REFERENCES (<privilege column list>) for all columns inherited from *T* in each of the subtables of *T*.

## Access Rules

*None.*

## General Rules

- 1) Case:
  - a) If GRANTED BY is omitted, then let *G* be OMITTED.
  - b) Otherwise, let *G* be <grantor>.
- 2) Let *A* be the result of applying the General Rules of [Subclause 12.8, “Grantor determination”](#), with *G* as *GRANTOR*.
- 3) Case:
  - a) If the <revoke statement> is a <revoke privilege statement>, and the applicable privileges for *A* do not include a privilege identifying *O*, then an exception condition is raised: *invalid grantor*.
  - b) If the <revoke statement> is a <revoke role statement>, then, for every role *R* identified by a <role revoked>, if the applicable roles of *A* do not include a role *AR* such that there exists a grantable role authorization descriptor with role *R* and grantee *A*, then an exception condition is raised: *invalid grantor*.
- 4) Case:
  - a) If the <revoke statement> is a <revoke privilege statement>, then, for every <grantee> specified, a set of privilege descriptors is identified. A privilege descriptor *P* is said to be *identified* if it belongs to the set of privilege descriptors that defined, for any <action> explicitly or implicitly in <privileges>, that <action> on *O*, or any of the objects in *S*, granted by *A* to <grantee>.  
NOTE 362 — Column privilege descriptors become identified when <action> explicitly or implicitly contains a <privilege column list>. Table/method descriptors become identified when <action> explicitly or implicitly contains a <privilege method list>.
  - b) If the <revoke statement> is a <revoke role statement>, then, for every <grantee> specified, a set of role authorization descriptors is identified. A role authorization descriptor is said to be *identified* if it defines the grant of any of the specified <role revoked>s to <grantee> with grantor *A*.
- 5) A privilege descriptor *D* is said to be *directly dependent* on another privilege descriptor *P* if  
Case:
  - a) If *D* identifies a view or a column of a view and the action of *D* is INSERT, UPDATE, or DELETE, then *D* is *directly dependent* on *P* if there exists a view privilege dependency descriptor whose supporting privilege descriptor is *P* and whose dependent privilege descriptor is *D*.

- b) Otherwise, *D* is directly dependent on *P* if one of the following is true:
- i) All of the following conditions hold:
    - 1) *P* indicates that the privilege that it represents is grantable.
    - 2) The grantee of *P* is the same as the grantor of *D*, or the grantee of *P* is PUBLIC, or, if the grantor of *D* is a <role name>, the grantee of *P* is an applicable role for the grantor of *D*.
    - 3) Case:
      - A) *P* and *D* are both column privilege descriptors. The action and the identified column of *P* are the same as the action and identified column of *D*, respectively.
      - B) *P* and *D* are both table privilege descriptors. The action and the identified table of *P* are the same as the action and the identified table of *D*, respectively.
      - C) *P* and *D* are both execute privilege descriptors. The action and the identified SQL-invoked routine of *P* are the same as the action and the identified SQL-invoked routine of *D*, respectively.
      - D) 13 *P* and *D* are both usage privilege descriptors. The action and the identified domain, character set, collation, transliteration, user-defined type, or sequence generator of *P* are the same as the action and the identified domain, character set, collation, transliteration, user-defined type, or sequence generator of *D*, respectively.
      - E) *P* and *D* are both under privilege descriptors. The action and the identified user-defined type or table of *P* are the same as the action and the identified user-defined type or table of *D*, respectively.
      - F) *P* and *D* are both table/method privilege descriptors. The action and the identified method and table of *P* are the same as the action and the identified method and table of *D*, respectively.
  - ii) All of the following conditions hold:
    - 1) The privilege descriptor for *D* indicates that its grantor is the special grantor value “\_SYSTEM”.
    - 2) The action of *P* is the same as the action of *D*.
    - 3) The grantee of *P* is the owner of the table, collation, or transliteration identified by *D* or the grantee of *P* is PUBLIC.
    - 4) One of the following conditions hold:
      - A) *P* and *D* are both column privilege descriptors, the privilege descriptor *D* identifies a <column name> *CVN* explicitly or implicitly contained in the <view column list> of a <view definition> *V*, and one of the following is true:
        - I) There exists a table *T* identified by a <table reference> contained in the <query expression> of *V* and a column *CT* that is a column of *T* and an underlying column of *CV*, such that the action for *P* is REFERENCES and either the identified column of *P* is *CT* or the identified table of *P* is *T*.
        - II) There exists a table *T* identified by a <table reference> contained in the <query expression> of *V* and a column *CT* that is a column of *T* and an underlying column

of *CV*, Such that the action for *P* is SELECT and either the identified column of *P* is *CT* or the identified table of *P* is *T*.

- B) The privilege descriptor *D* identifies the <collation name> of a <collation definition> *CO* and the identified character set name of *P* is included in the collation descriptor for *CO*, or the identified transliteration name of *P* is included in the collation descriptor for *CO*.
- C) The privilege descriptor *D* identifies the <transliteration name> of a <transliteration definition> *TD* and the identified character set name of *P* is contained in the <source character set specification> or the <target character set specification> immediately contained in *TD*.

iii) All of the following conditions hold:

- 1) The privilege descriptor for *D* indicates that its grantor is the special grantor value “\_SYSTEM”.
- 2) The grantee of *P* is the owner of the domain identified by *D* or the grantee of *P* is PUBLIC.
- 3) The privilege descriptor *D* identifies the <domain name> of a <domain definition> *DO* and either the column privilege descriptor *P* has an action of REFERENCES and identifies a column referenced in the <search condition> included in the domain descriptor for *DO*, or the privilege descriptor *P* has an action of USAGE and identifies a domain, collation, character set, or transliteration whose <domain name>, <collation name>, <character set name>, or <transliteration name>, respectively, is contained in the <search condition> of the domain descriptor for *DO*.

6) The *privilege dependency graph* is a directed graph such that all of the following are true:

- a) Each node represents a privilege descriptor.
- b) Each arc from node *P1* to node *P2* represents the fact that *P2* directly depends on *P1*.

An *independent node* is a node that has no incoming arcs.

7) A privilege descriptor *P* is said to be *modified* if all of the following are true:

- a) *P* indicates that the privilege that it represents is grantable.
- b) *P* directly depends on an identified privilege descriptor or a modified privilege descriptor.
- c) Case:
  - i) If *P* is neither a SELECT nor a REFERENCES column privilege descriptor that identifies a <column name> *CVN* explicitly or implicitly contained in the <view column list> of a <view definition> *V*, then let *XO* and *XA* respectively be the identifier of the object identified by a privilege descriptor *X* and the action of *X*. Within the set of privilege descriptors upon which *P* directly depends, there exist some *XO* and *XA* for which the set of identified privilege descriptors unioned with the set of modified privilege descriptors include all privilege descriptors specifying the grant of *XA* on *XO* WITH GRANT OPTION.
  - ii) If *P* is a column privilege descriptor that identifies a column *CV* identified by a <column name> *CVN* explicitly or implicitly contained in the <view column list> of a <view definition> *V* with an action *PA* of REFERENCES or SELECT, then let *SP* be the set of privileges upon which *P* directly depends. For every table *T* identified by a <table reference> contained in the <query

expression> of  $V$ , let  $RT$  be the <table name> of  $T$ . There exists a column  $CT$  whose <column name> is  $CRT$ , such that all of the following are true:

- 1)  $CT$  is a column of  $T$  and an underlying column of  $CV$ .
  - 2) Every privilege descriptor  $PD$  that is the descriptor of some member of  $SP$  that specifies the action  $PA$  on  $CRT$  WITH GRANT OPTION is either an identified privilege descriptor for  $CRT$  or a modified privilege descriptor for  $CRT$ .
- d) At least one of the following is true:
- i) GRANT OPTION FOR is specified and the grantor of  $P$  is the special grantor value “\_SYSTEM”.
  - ii) There exists a path to  $P$  from an independent node that includes no identified or modified privilege descriptors.  $P$  is said to be a *marked modified privilege descriptor*.
  - iii)  $P$  directly depends on a marked modified privilege descriptor, and the grantor of  $P$  is the special grantor value “\_SYSTEM”.  $P$  is said to be a *marked modified privilege descriptor*.
- 8) A role authorization descriptor  $D$  is said to be *directly dependent* on another role authorization descriptor  $RD$  if all of the following conditions hold:
- a)  $RD$  indicates that the role that it represents is grantable.
  - b) The role name of  $D$  is the same as the role name of  $RD$ .
  - c) The grantee of  $RD$  is the same as the grantor of  $D$ , or the grantee of  $RD$  is PUBLIC, or, if the grantor of  $D$  is a <role name>, the grantee of  $RD$  is an applicable role for the grantor of  $D$ .
- 9) The *role dependency graph* is a directed graph such that all of the following are true:
- a) Each node represents a role authorization descriptor.
  - b) Each arc from node  $R1$  to node  $R2$  represents the fact that  $R2$  directly depends on  $R1$ .

An independent node is one that has no incoming arcs.

- 10) A role authorization descriptor  $RD$  is said to be *abandoned* if it is not an independent node, and it is not itself an identified role authorization descriptor, and there exists no path to  $RD$  from any independent node other than paths that include an identified role authorization descriptor.
- 11) An arc from a node  $P$  to a node  $D$  of the privilege dependency graph is said to be *unsupported* if all of the following are true:
- a) The grantor of  $D$  and the grantee of  $P$  are both <role name>s.
  - b) The destruction of all abandoned role authorization descriptors and, if ADMIN OPTION FOR is not specified, all identified role authorization descriptors would result in the grantee of  $P$  no longer being an applicable role for the grantor of  $D$ .

- 12) A privilege descriptor  $P$  is *abandoned* if

Case:

- a) It is not an independent node, and  $P$  is not itself an identified or a modified privilege descriptor, and there exists no path to  $P$  from any independent node other than paths that include an identified privilege descriptor or a modified privilege descriptor or an unsupported arc, and, if <revoke statement> specifies WITH HIERARCHY OPTION, then  $P$  has the WITH HIERARCHY OPTION.



- b) All of the following conditions hold:
  - i)  $P$  is a column privilege descriptor that identifies a <column name>  $CVN$  explicitly or implicitly contained in the <view column list> of a <view definition>  $V$ , with an action  $PA$  of REFERENCES or SELECT.
  - ii) Letting  $SP$  be the set of privileges upon which  $P$  directly depends, at least one of the following is true:
    - 1) There exists some table name  $RT$  such that all of the following are true:
      - A)  $RT$  is the name of the table identified by some <table reference> contained in the <query expression> of  $V$ .
      - B) For every column privilege descriptor  $CPD$  that is the descriptor of some member of  $SP$  that specifies the action  $PA$  on  $RT$ ,  $CPD$  is either an identified privilege descriptor for  $RT$  or an abandoned privilege descriptor for  $RT$ .
    - 2) There exists some column name  $CRT$  such that all of the following are true:
      - A)  $CRT$  is the name of some column of the table identified by some <table reference> contained in the <query expression> of  $V$ .
      - B) For every column privilege descriptor  $CPD$  that is the descriptor of some member of  $SP$  that specifies the action  $PA$  on  $CRT$ ,  $CPD$  is either an identified privilege descriptor for  $CRT$  or an abandoned privilege descriptor for  $CRT$ .

13) The *revoke destruction action* is defined as

Case:

- a) If the <revoke statement> is a <revoke privilege statement>, then

Case:

- i) If the <revoke statement> specifies the WITH HIERARCHY OPTION, then the removal of the WITH HIERARCHY OPTION from all identified and abandoned privilege descriptors.
  - ii) Otherwise, the destruction of all abandoned privilege descriptors and, if GRANT OPTION FOR is not specified, all identified privilege descriptors.
- b) If the <revoke statement> is a <revoke role statement>, then the destruction of all abandoned role authorization descriptors, all abandoned privilege descriptors and, if GRANT ADMIN OPTION FOR is not specified, all identified role authorization descriptors.

14) Let  $SI$  be the name of any schema and  $AI$  be the <authorization identifier> that owns the schema identified by  $SI$ .

15) Let  $V$  be any view descriptor included in  $SI$ . Let  $QE$  be the <query expression> of  $V$ .  $V$  is said to be *abandoned* if the revoke destruction action would result in  $AI$  no longer having in its applicable privileges any of the following:

- a) SELECT privilege on at least one column of every table identified by a <table reference> is contained in  $QE$ .
- b) SELECT privilege on every column identified by a <column reference> contained in  $QE$ .

- c) USAGE privilege on every domain, every collation, every character set, and every transliteration whose names are contained in *QE*.
  - d) USAGE privilege on any user-defined type *UDT* such that some <data type> contained in *V* is usage-dependent on *UDT*.
  - e) 04 EXECUTE privilege on every SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in *QE*.
  - f) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in *QE* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is subject routine of *MR*.
  - g) SELECT privilege on any column identified by a <column reference> contained in the <scalar subquery> that is equivalent to some <dereference operation> contained in *QE*.
  - h) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of any <reference resolution> that is contained in *QE*.
  - i) SELECT privilege on the scoped table of any <reference resolution> that is contained in *QE*.
  - j) If *V* is the descriptor of a referenceable table, then USAGE privilege on the structured type associated with the view described by *V*.
  - k) UNDER privilege on every direct supertable of the view described by *V*.
  - l) SELECT privilege WITH HIERARCHY OPTION privilege on at least one supertable of every typed table identified by a <table reference> that simply contains an <only spec> and that is contained in *QE*.
- 16) 09 Let *T* be any table descriptor included in *SI*. *T* is said to be *abandoned* if the revoke destruction action would result in *AI* no longer having any of the following:
- a) If *T* is the descriptor of a referenceable table, then USAGE privilege on the structured type associated with the table described by *T*.
  - b) UNDER privilege on every direct supertable of the table described by *T*.
- 17) Let *TC* be any table constraint descriptor included in *SI*. *TC* is said to be *abandoned* if the revoke destruction action would result in *AI* no longer having in its applicable privileges any of the following:
- a) REFERENCES privilege on at least one column of every table identified by a <table reference> contained in the applicable <search condition> of *TC*.
  - b) REFERENCES privilege on every column identified by a <column reference> contained in the applicable <search condition> of *TC*.
  - c) USAGE privilege on every domain, every collation, every character set, and every transliteration whose names are contained in any applicable <search condition> of *TC*.
  - d) USAGE privilege on any user-defined type *UDT* such that some <data type> contained in the applicable <search condition> of *TC* is usage-dependent on *UDT*.
  - e) 04 EXECUTE privilege on every SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any applicable <search condition> of *TC*.

- f) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in any applicable <search condition> of *TC* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the subject routine of *MR*.
  - g) SELECT privilege on any column identified by a <column reference> contained in the <scalar subquery> that is equivalent to some <dereference operation> contained in any applicable <search condition> of *TC*.
  - h) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of any <reference resolution> that is contained in any applicable <search condition> of *TC*.
  - i) SELECT privilege on the scoped table of any <reference resolution> that is contained in any applicable <search condition> of *TC*.
  - j) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of every typed table identified by a <table reference> that simply contains an <only spec> and that is contained in the applicable <search condition> of *TC*.
- 18) Let *AX* be any assertion descriptor included in *SI*. *AX* is said to be *abandoned* if the revoke destruction action would result in *AI* no longer having in its applicable privileges any of the following:
- a) REFERENCES privilege on at least one column of every table identified by a <table reference> contained in the applicable <search condition> of *AX*.
  - b) REFERENCES privilege on every column identified by a <column reference> contained in the applicable <search condition> of *AX*.
  - c) USAGE privilege on every domain, every collation, every character set, and every transliteration whose names are contained in any applicable <search condition> of *AX*.
  - d) USAGE privilege on any user-defined type *UDT* such that some <data type> contained in the applicable <search condition> of *AX* is usage-dependent on *UDT*.
  - e) 04 EXECUTE privilege on every SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any applicable <search condition> of *AX*.
  - f) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in the applicable <search condition> of *AX* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the subject routine of *MR*.
  - g) SELECT privilege on any column identified by a <column reference> contained in the <scalar subquery> that is equivalent to some <dereference operation> contained in any applicable <search condition> of *AX*.
  - h) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of any <reference resolution> that is contained in any applicable <search condition> of *AX*.
  - i) SELECT privilege on the scoped table of any <reference resolution> that is contained in any applicable <search condition> of *AX*.
  - j) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of every typed table identified by a <table reference> that simply contains an <only spec> and that is contained in the applicable <search condition> of *AX*.
- 19) 04 Let *TR* be any trigger descriptor included in *SI*. *TR* is said to be *abandoned* if the revoke destruction action would result in *AI* no longer having in its applicable privileges any of the following:

- a) TRIGGER privilege on the subject table of *TR*.
- b) REFERENCES privilege on at least one column of every table identified by a <table reference> contained in any <search condition> of *TR*.
- c) SELECT privilege on every column identified by a <column reference> contained in any <search condition> of *TR*.
- d) USAGE privilege on every domain, collation, character set, and transliteration whose name is contained in any <search condition> of *TR*.
- e) USAGE privilege on any user-defined type *UDT* such that some <data type> contained in any <search condition> of *TR* is usage-dependent on *UDT*.
- f) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in any <search condition> of *TR* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the subject routine of *MR*.
- g) EXECUTE privilege on the SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any <search condition> of *TR*.
- h) EXECUTE privilege on the SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in the <triggered SQL statement> of *TR*.
- i) SELECT privilege on at least one column of every table identified by a <table reference> contained in a <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
- j) SELECT privilege on at least one column of every table identified by a <table reference> contained in a <table expression> or <select list> immediately contained in a <select statement: single row> contained in the <triggered SQL statement> of *TR*.
- k) SELECT privilege on at least one column of every table identified by a <table reference> and <column reference> contained in a <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
- l) SELECT privilege on at least one column of every table identified by a <table reference> and <column reference> contained in a <value expression> simply contained in an <update source> or an <assigned row> contained in the <triggered SQL statement> of *TR*.
- m) INSERT privilege on every column

Case:

- i) Identified by a <column name> contained in the <insert column list> of an <insert statement> or a <merge statement> contained in the <triggered SQL statement> of *TR*.
- ii) Of the table identified by the <table name> immediately contained in an <insert statement> that does not contain an <insert column list> and that is contained in the <triggered SQL statement> of *TR*.
- iii) Of the table identified by the <target table> contained in a <merge statement> that contains a <merge insert specification> and that does not contain an <insert column list> and that is contained in the <triggered SQL statement> of *TR*.

- n) UPDATE privilege on every column identified by a <column name> is contained in an <object column> contained in either an <update statement: positioned>, an <update statement: searched>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
- o) DELETE privilege on every table identified by a <table name> contained in either a <delete statement: positioned> or a <delete statement: searched> contained in the <triggered SQL statement> of *TR*.
- p) USAGE privilege on every domain, collation, character set, transliteration, and sequence generator whose name is contained in the <triggered SQL statement> of *TR*.
- q) USAGE privilege on any user-defined type *UDT* such that some <data type> contained in the <triggered SQL statement> of *TR* is usage-dependent on *UDT*.
- r) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in any <triggered SQL statement> of *TR* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the subject routine of *MR*.
- s) SELECT privilege on any column identified by a <column reference> contained in the <scalar subquery> that is equivalent to some <dereference operation> contained in any of the following:
  - i) A <search condition> of *TR*.
  - ii) A <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
  - iii) A <table expression> or <select list> immediately contained in a <select statement: single row> contained in the <triggered SQL statement> of *TR*.
  - iv) A <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
  - v) A <value expression> contained in an <update source> or an <assigned row> contained in the <triggered SQL statement> of *TR*.
- t) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of any <reference resolution> that is contained in any of the following:
  - i) A <search condition> of *TR*.
  - ii) A <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
  - iii) A <table expression> or <select list> immediately contained in a <select statement: single row> contained in the <triggered SQL statement> of *TR*.
  - iv) A <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
  - v) A <value expression> contained in an <update source> or an <assigned row> contained in the <triggered SQL statement> of *TR*.
- u) SELECT privilege on the scoped table of any <reference resolution> contained in any of the following:
  - i) A <search condition> of *TR*.
  - ii) A <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.

- iii) A <table expression> or <select list> immediately contained in a <select statement: single row> contained in the <triggered SQL statement> of *TR*.
  - iv) A <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <triggered SQL statement> of *TR*.
  - v) A <value expression> contained in an <update source> or an <assigned row> contained in the <triggered SQL statement> of *TR*.
  - v) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of every typed table identified by a <table reference> that simply contains an <only spec> and that is contained in the <triggered SQL statement> of *TR*.
- 20) Let *DC* be any domain constraint descriptor included in *SI*. *DC* is said to be *abandoned* if the revoke destruction action would result in *AI* no longer having in its applicable privileges any of the following:
- a) REFERENCES privilege on at least one column of every table identified by a <table reference> contained in *TR*.
  - b) REFERENCES privilege on every column identified by a <column reference> contained in the <search condition> of *DC*.
  - c) USAGE privilege on every domain, every user-defined type, every collation, every character set, and every transliteration whose names are contained in any <search condition> of *DC*.
  - d) USAGE privilege on any user-defined type *UDT* such that some <data type> contained in any <search condition> of *DC* is usage-dependent on *UDT*.
  - e) 04 EXECUTE privilege on every SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in any <search condition> of *DC*.
  - f) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in any <search condition> of *DC* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the subject routine of *MR*.
  - g) SELECT privilege on any column identified by a <column reference> contained in a <scalar subquery> that is equivalent to some <dereference operation> contained in any <search condition> of *DC*.
  - h) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of any <reference resolution> that is contained in any <search condition> of *DC*.
  - i) SELECT privilege on the scoped table of any <reference resolution> that is contained in contained in any <search condition> of *DC*.
  - j) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of every typed table identified by a <table reference> that simply contains an <only spec> and that is contained in the <triggered SQL statement> of *TR*.
- 21) For every domain descriptor *DO* included in *SI*, *DO* is said to be *lost* if the revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on every character set included in the data type descriptor included in *DO*.
- 22) For every table descriptor *TD* contained in *SI*, for every column descriptor *CD* included in *TD*, *CD* is said to be *lost* if any of the following are true:

- a) The revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on any character set included in the data type descriptor included in *CD*.
  - b) The revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on any user-defined type *UDT* such that a data type descriptor included in *CD* describes a type that is usage-dependent on *UDT*.
  - c) The name of the domain *DN* included in *CD*, if any, identifies a lost domain descriptor and the revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on any character set included in the data type descriptor of the domain descriptor of *DN*.
- 23) For every SQL-client module *MO*, let *G* be the <module authorization identifier> that owns *MO*. *MO* is said to be *lost* if the revoke destruction action would result in *G* no longer having in its applicable privileges USAGE privilege on the character set referenced in the <module character set specification> of *MO*.
- 24) For every user-defined type descriptor *DT* included in *SI*, *DT* is said to be *abandoned* if any of the following are true:
- a) The revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on any user-defined type *UDT* such that a data type descriptor included in *DT* describes a type that is usage-dependent on *UDT*.
  - b) 13 The revoke destruction action would result in *AI* no longer having in its applicable privileges the UNDER privilege on any user-defined type that is a direct supertype of *DT*.
- 25) *SI* is said to be *lost* if the revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on the default character set included in the *SI*.
- 26) For every collation descriptor *CN* contained in *SI*, *CN* is said to be *impacted* if the revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on the collation whose name is contained in the <existing collation name> of *CN*.
- 27) For every character set descriptor *CSD* contained in *SI*, *CSD* is said to be *impacted* if the revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on the collation whose name is contained in *CSD*.
- 28) 13 For every descriptor included in *SI* that includes a data type descriptor *DTD*, *DTD* is said to be *impacted* if the revoke destruction action would result in *AI* no longer having in its applicable privileges USAGE privilege on the collation whose name is included in *DTD*.
- 29) Let *RD* be any routine descriptor with an SQL security characteristic of DEFINER that is included in *SI*. *RD* is said to be *abandoned* if the revoke destruction action would result in *AI* no longer having in its applicable privileges all of the following:
- a) 04 EXECUTE privilege on the SQL-invoked routine that is the subject routine of any <routine invocation>, <method invocation>, <static method invocation>, or <method reference> that is contained in the <routine body> of *RD*.
  - b) SELECT privilege on at least one column of each table identified by a <table reference> contained in a <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the SQL routine body of *RD*.
  - c) SELECT privilege on at least one column of each table identified by a <table reference> contained in a <table expression> or <select list> immediately contained in a <select statement: single row> contained in the SQL routine body of *RD*.

- d) SELECT privilege on at least one column of each table identified by a <table reference> contained in a <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the SQL routine body of *RD*.
- e) SELECT privilege on at least one column of each table identified by a <table reference> contained in a <value expression> simply contained in an <update source> or an <assigned row> contained in the SQL routine body of *RD*.
- f) SELECT privilege on at least one column identified by a <column reference> contained in a <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <SQL routine body> of *RD*.
- g) SELECT privilege on at least one column identified by a <column reference> contained in a <value expression> simply contained in an <update source> or an <assigned row> contained in the SQL routine body of *RD*.
- h) INSERT privilege on each column

Case:

- i) Identified by a <column name> contained in the <insert column list> of an <insert statement> or a <merge statement> contained in the SQL routine body of *RD*.
- ii) Of the table identified by the <table name> immediately contained in an <insert statement> that does not contain an <insert column list> and that is contained in the SQL routine body of *RD*.
- iii) Of the table identified by the <target table> immediately contained in a <merge statement> that contains a <merge insert specification> and that does not contain an <insert column list> and that is contained in the SQL routine body of *RD*.
- i) UPDATE privilege on each column whose name is contained in an <object column> contained in either an <update statement: positioned>, an <update statement: searched>, or a <merge statement> contained in the SQL routine body of *RD*.
- j) DELETE privilege on each table whose name is contained in a <table name> contained in either a <delete statement: positioned> or a <delete statement: searched> contained in the SQL routine body of *RD*.
- k) USAGE privilege on each domain, collation, character set, transliteration, and sequence generator whose name is contained in the SQL routine body of *RD*.
- l) USAGE privilege on each user-defined type *UDT* such that a declared type of any SQL parameter, returns data type, or result cast included in *RD* is usage-dependent on *UDT*.
- m) USAGE privilege on each user-defined type *UDT* such that some <data type> contained in the SQL routine body of *RD* is usage-dependent on *UDT*.
- n) The table/method privilege on every table *TI* and every method *M* such that there is a <method reference> *MR* contained in the SQL routine body of *RI* such that *TI* is in the scope of the <value expression primary> of *MR* and *M* is the subject routine of *MR*.
- o) SELECT privilege on any column identified by a <column reference> contained in a <scalar subquery> that is equivalent to a <dereference operation> contained in any of the following:
  - i) A <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the <SQL routine body> of *RD*.



- ii) A <table expression> or <select list> immediately contained in a <select statement: single row> contained in the <SQL routine body> of *RD*.
  - iii) A <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <SQL routine body> of *RD*.
  - iv) A <value expression> contained in an <update source> or an <assigned row> contained in the <SQL routine body> of *RD*.
- p) SELECT privilege WITH HIERARCHY OPTION on at least one supertable of the scoped table of any <reference resolution> that is contained in any of the following:
- i) A <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the SQL routine body of *RD*.
  - ii) A <table expression> or <select list> immediately contained in a <select statement: single row> contained in the SQL routine body of *RD*.
  - iii) A <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the SQL routine body of *RD*.
  - iv) A <value expression> simply contained in an <update source> or an <assigned row> contained in the SQL routine body of *RD*.
- q) SELECT privilege on the scoped table of any <reference resolution> that is contained in any of the following:
- i) A <query expression> simply contained in a <cursor specification>, an <insert statement>, or a <merge statement> contained in the <SQL routine body> of *RD*.
  - ii) A <table expression> or <select list> immediately contained in a <select statement: single row> contained in the <SQL routine body> of *RD*.
  - iii) A <search condition> contained in a <delete statement: searched>, an <update statement: searched>, or a <merge statement> contained in the <SQL routine body> of *RD*.
  - iv) A <value expression> contained in an <update source> or an <assigned row> contained in the <SQL routine body> of *RD*.
- r) 13 SELECT privilege WITH HIERARCHY OPTION on at least one supertable of every typed table identified by a <table reference> that simply contains an <only spec> and that is contained in the <SQL routine body> of *RD*.
- 30) 09 For every table descriptor *TD* included in *SI*, for every column descriptor *CD* included in *TD*, *CD* is said to be *contaminated* if *CD* includes one of the following:
- a) A user-defined type descriptor that describes a supertype of a user-defined type described by an abandoned user-defined type descriptor.
  - b) A reference type descriptor that includes a user-defined type descriptor that describes a supertype of a user-defined type described by an abandoned user-defined type descriptor.
  - c) A collection type descriptor that includes a user-defined type descriptor that describes a supertype of a user-defined type described by an abandoned user-defined type descriptor.

- d) A collection type descriptor that includes a reference type descriptor that includes a user-defined type descriptor that describes a supertype of a user-defined type described by an abandoned user-defined type descriptor.
- 31) 09 13 If RESTRICT is specified, and there exists an abandoned privilege descriptor, abandoned view, abandoned table constraint, abandoned assertion, abandoned domain constraint, lost domain, lost column, lost schema, or a descriptor that includes an impacted data type descriptor, impacted collation, impacted character set, abandoned user-defined type, forsaken column descriptor, forsaken domain descriptor, or abandoned routine descriptor, then an exception condition is raised: *dependent privilege descriptors still exist*.
- 32) 04 If CASCADE is specified, then the impact on an SQL-client module that is determined to be a lost module is implementation-defined.
- 33) Case:
- a) If the <revoke statement> is a <revoke privilege statement>, then
 

Case:

    - i) If neither WITH HIERARCHY OPTION nor GRANT OPTION FOR is specified, then:
      - 1) All abandoned privilege descriptors are destroyed.
      - 2) The identified privilege descriptors are destroyed.
      - 3) The modified privilege descriptors are set to indicate that they are not grantable.
    - ii) If WITH HIERARCHY OPTION is specified, then the WITH HIERARCHY OPTION is removed from all identified and abandoned privilege descriptors, if present.
    - iii) If GRANT OPTION FOR is specified, then
 

Case:

      - 1) If CASCADE is specified, then all abandoned privilege descriptors are destroyed.
      - 2) Otherwise, if there are any privilege descriptors directly dependent on an identified privilege descriptor that are not modified privilege descriptors, then an exception condition is raised: *dependent privilege descriptors still exist*.

The identified privilege descriptors and the modified privilege descriptors are set to indicate that they are not grantable.
  - b) If the <revoke statement> is a <revoke role statement>, then:
    - i) If CASCADE is specified, then all abandoned role authorization descriptors are destroyed.
    - ii) All abandoned privilege descriptors are destroyed.
    - iii) Case:
      - 1) If ADMIN OPTION FOR is specified, then the identified role authorization descriptors are set to indicate that they are not grantable.
      - 2) If ADMIN OPTION FOR is not specified, then the identified role authorization descriptors are destroyed.

- 34) For every abandoned view descriptor *V*, let *S1.VN* be the <table name> of *V*. The following <drop view statement> is effectively executed without further Access Rule checking:

```
DROP VIEW S1.VN CASCADE
```

- 35) For every abandoned table descriptor *T*, let *S1.TN* be the <table name> of *T*. The following <drop table statement> is effectively executed without further Access Rule checking:

```
DROP TABLE S1.TN CASCADE
```

- 36) For every abandoned table constraint descriptor *TC*, let *S1.TCN* be the <constraint name> of *TC* and let *S2.T2* be the <table name> of the table that contains *TC* (*S1* and *S2* possibly equivalent). The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE S2.T2
 DROP CONSTRAINT S1.TCN
 CASCADE
```

- 37) For every abandoned assertion descriptor *AX*, let *S1.AXN* be the <constraint name> of *AX*. The following <drop assertion statement> is effectively executed without further Access Rule checking:

```
DROP ASSERTION S1.AXN CASCADE
```

- 38) For every abandoned trigger descriptor *TR*, let *S1.TRN* be the <trigger name> of *TR*. The following <drop trigger statement> is effectively executed without further Access Rule checking:

```
DROP TRIGGER S1.TRN
```

- 39) For every abandoned domain constraint descriptor *DC*, let *S1.DCN* be the <constraint name> of *DC* and let *S2.DN* be the <domain name> of the domain that contains *DC*. The following <alter domain statement> is effectively executed without further Access Rule checking:

```
ALTER DOMAIN S2.DN
 DROP CONSTRAINT S1.DCN
```

- 40) For every lost column descriptor *CD*, let *S1.TN* be the <table name> of the table whose descriptor includes the descriptor *CD* and let *CN* be the <column name> of *CD*. The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE S1.TN
 DROP COLUMN CN CASCADE
```

- 41) For every lost domain descriptor *DO*, let *S1.DN* be the <domain name> of *DO*. The following <drop domain statement> is effectively executed without further Access Rule checking:

```
DROP DOMAIN S1.DN CASCADE
```

- 42) For every lost schema *S1*, the default character set of that schema is modified to include the name of the implementation-defined <character set specification> that would have been this schema's default character set had the <schema definition> not specified a <schema character set specification>.

- 43) If the object identified by *O* is a collation, let *OCN* be the name of that collation.

## 12.7 &lt;revoke statement&gt;

- 44) For every descriptor that includes an impacted data type descriptor *DTD*, *DTD* is modified such that it does not include *OCN*.
- 45) For every impacted collation descriptor *CD* with included collation name *CN*, the following <drop collation statement> is effectively executed without further Access Rule checking:

```
DROP COLLATION CN CASCADE
```

- 46) For every impacted character set descriptor *CSD* with included character set name *CSN*, *CSD* is modified so that the included collation name is the name of the default collation for the character set on which *CSD* is based.
- 47) For every abandoned user-defined type descriptor *DT* with <user-defined type name> *SI.DTN*, the following <drop data type statement> is effectively executed without further Access Rule checking:

```
DROP TYPE SI.DTN CASCADE
```

- 48) 04 13 For every abandoned SQL-invoked routine descriptor *RD*, let *R* be the SQL-invoked routine whose descriptor is *RD*. Let *SN* be the <specific name> of *R*. The following <drop routine statement> is effectively executed without further Access Rule checking:

```
DROP SPECIFIC ROUTINE SN CASCADE
```

- 49) If the <revoke statement> is a <revoke privilege statement>, then:
- a) For every combination of <grantee> and <action> on *O* specified in <privileges>, if there is no corresponding privilege descriptor in the set of identified privilege descriptors, then a completion condition is raised: *warning — privilege not revoked*.
  - b) If ALL PRIVILEGES was specified, then for each <grantee>, if no privilege descriptors were identified, then a completion condition is raised: *warning — privilege not revoked*.
- 50) For every contaminated column descriptor *CD*, let *SI.TN* be the <table name> of the table whose descriptor includes the descriptor *CD* and let *CN* be the <column name> of *CD*. The following <alter table statement> is effectively executed without further Access Rule checking:

```
ALTER TABLE SI.TN
 DROP COLUMN CN CASCADE
```

## Conformance Rules

- 1) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <revoke role statement>.
- 2) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke statement> that contains a <drop behavior> that contains CASCADE.
- 3) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke option extension> that contains GRANT OPTION FOR.
- 4) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke statement> that contains a <privileges> that contains an <object name> where the owner of the SQL-schema that is specified explicitly or implicitly in the <object name> is not the current authorization identifier.

- 5) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke statement> such that there exists a privilege descriptor *PD* that satisfies all the following conditions:
- a) *PD* identifies the object identified by <object name> simply contained in <privileges> contained in the <revoke statement>.
  - b) *PD* identifies the <grantee> identified by any <grantee> simply contained in <revoke statement> and that <grantee> does not identify the owner of the SQL-schema that is specified explicitly or implicitly in the <object name> simply contained in <privileges> contained in the <revoke statement>.
  - c) *PD* identifies the action identified by the <action> simply contained in <privileges> contained in the <revoke statement>.
  - d) *PD* indicates that the privilege is grantable.
- 6) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <revoke option extension> that contains HIERARCHY OPTION FOR.

## 12.8 Grantor determination

### Function

Determine the grantor of a privilege or role authorization, or the intended owner of a role.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *G* be the *GRANTOR* specified in an application of this Subclause.
- 2) The grantor *A* is derived from *G* as follows.

Case:

- a) If *G* is OMITTED, then

Case:

- i) If there is a current user identifier, then *A* is the current user identifier.
- ii) Otherwise, *A* is the current role name.

- b) If *G* is CURRENT\_USER, then

Case:

- i) If there is no current user identifier, then an exception condition is raised: *invalid grantor*.
- ii) Otherwise, *A* is the current user identifier.

- c) If *G* is CURRENT\_ROLE, then

Case:

- i) If there is no current role name, then an exception condition is raised: *invalid grantor*.
- ii) Otherwise, *A* is the current role name.

### Conformance Rules

- 1) Without Feature T332, “Extended roles”, conforming SQL language shall contain no <grantor>.

## 13 SQL-client modules

### 13.1 <SQL-client module definition>

*This Subclause is modified by Subclause 14.1, “<SQL-client module definition>”, in ISO/IEC 9075-9.*

#### Function

Define an SQL-client module.

#### Format

```
<SQL-client module definition> ::=
 <module name clause> <language clause> <module authorization clause>
 [<module path specification>]
 [<module transform group specification>]
 [<module collations>]
 [<temporary table declaration>...]
 <module contents>...

<module authorization clause> ::=
 SCHEMA <schema name>
 | AUTHORIZATION <module authorization identifier>
 [FOR STATIC { ONLY | AND DYNAMIC }]
 | SCHEMA <schema name> AUTHORIZATION <module authorization identifier>
 [FOR STATIC { ONLY | AND DYNAMIC }]

<module authorization identifier> ::=
 <authorization identifier>

<module path specification> ::=
 <path specification>

<module transform group specification> ::=
 <transform group specification>

<module collations> ::=
 <module collation specification>...

<module collation specification> ::=
 COLLATION <collation name> [FOR <character set specification list>]

<character set specification list> ::=
 <character set specification> [{ <comma> <character set specification> }...]

<module contents> ::=
 <declare cursor>
 | <dynamic declare cursor>
 | <externally-invoked procedure>
```

## Syntax Rules

- 1) The <language clause> shall not specify SQL.
- 2) If SCHEMA <schema name> is not specified, then a <schema name> equivalent to <module authorization identifier> is implicit.
- 3) If the explicit or implicit <schema name> does not specify a <catalog name>, then an implementation-defined <catalog name> is implicit.
- 4) The implicit or explicit <catalog name> is the implicit <catalog name> for all unqualified <schema name>s in the <SQL-client module definition>.
- 5) If <module path specification> is not specified, then a <module path specification> containing an implementation-defined <schema name list> that contains the <schema name> contained in <module authorization clause> is implicit.
- 6) The explicit or implicit <catalog name> of each <schema name> contained in the <schema name list> of the <module path specification> shall be equivalent to the <catalog name> of the explicit or implicit <schema name> contained in <module authorization clause>.
- 7) The <schema name list> of the explicit or implicit <module path specification> is used as the SQL-path of the <SQL-client module definition>. The SQL-path is used to effectively qualify unqualified <routine name>s that are immediately contained in <routine invocation>s that are contained in the <SQL-client module definition>.
- 8) Case:
  - a) If <module transform group specification> is not specified, then a <module transform group specification> containing a <multiple group specification> with a <group specification> *GS* for each <host parameter declaration> contained in <host parameter declaration list> of each <externally-invoked procedure> contained in <SQL-client module definition> whose <host parameter data type> *UDT* identifies a user-defined type with no <locator indication> is implicit. The <group name> of *GS* is implementation-defined and its <path-resolved user-defined type name> is *UDT*.
  - b) If <module transform group specification> contains a <single group specification> with a <group name> *GN*, then a <module transform group specification> containing a <multiple group specification> that contains a <group specification> *GS* for each <host parameter declaration> contained in <host parameter declaration list> of each <externally-invoked procedure> contained in <SQL-client module definition> whose <host parameter data type> *UDT* identifies a user-defined type with no <locator indication> is implicit. The <group name> of *GS* is *GN* and its <path-resolved user-defined type name> is *UDT*.
  - c) If <module transform group specification> contains a <multiple group specification> *MGS*, then a <module transform group specification> containing <multiple group specification> that contains *MGS* extended with a <group specification> *GS* for each <host parameter declaration> contained in <host parameter declaration list> of each <externally-invoked procedure> contained in <SQL-client module definition> whose <host parameter data type> *UDT* identifies a user-defined type with no <locator indication> and no equivalent of *UDT* is contained in any <group specification> contained in *MGS* is implicit. The <group name> of *GS* is implementation-defined and its <path-resolved user-defined type name> is *UDT*.
- 9) No two <character set specification>s contained in any <module collation specification> shall be equivalent.



- 10) A <module collation specification> *MCS* specifies the SQL-client module collation for one or more character sets for the SQL-client module. Let *CO* be the collation identified by the <collation name> contained in *MCS*.

Case:

- a) If <character set specification list> is specified, then the collation specified by *CO* shall be applicable to every character set identified by a <character set specification> simply contained in the <module collation specification>. For each character set specified, the SQL-client module collation for that character set is set to *CO*.
  - b) Otherwise, the character sets for which the SQL-client module collation is set to *CO* are implementation-defined.
- 11) A <declare cursor> shall precede in the text of the <SQL-client module definition> any <externally-invoked procedure> that references the <cursor name> of the <declare cursor>.
- 12) A <dynamic declare cursor> shall precede in the text of the <SQL-client module definition> any <externally-invoked procedure> that references the <cursor name> of the <dynamic declare cursor>.
- 13) If neither FOR STATIC ONLY nor FOR STATIC AND DYNAMIC is specified, then FOR STATIC AND DYNAMIC is implicit.
- 14) For every <declare cursor> in an <SQL-client module definition>, the <SQL-client module definition> shall contain exactly one <open statement> that specifies the <cursor name> declared in the <declare cursor>.

NOTE 363 — See the Syntax Rules of Subclause 14.1, “<declare cursor>”.

- 15) Let *EIP1* and *EIP2* be two <externally-invoked procedure>s contained in an <SQL-client module definition> that have the same number of <host parameter declaration>s and immediately contain a <fetch statement> referencing the same <cursor name>. Let *n* be the number of <host parameter declaration>s. Let *PI<sub>i</sub>*, 1 (one) ≤ *i* ≤ *n*, be the *i*-th <host parameter declaration> of *EIP1*. Let *DT1<sub>i</sub>* be the <data type> contained in *PI<sub>i</sub>*. Let *P2<sub>i</sub>* be the *i*-th <host parameter declaration> of *EIP2*. Let *DT2<sub>i</sub>* be the <data type> contained in *P2<sub>i</sub>*. For each *i*, 1 (one) ≤ *i* ≤ *n*,

Case:

- a) If *DT1<sub>i</sub>* and *DT2<sub>i</sub>* both identify a binary large object type, then either *PI<sub>i</sub>* and *P2<sub>i</sub>* shall both be binary large object locator parameters or neither shall be binary large object locator parameters.
- b) If *DT1<sub>i</sub>* and *DT2<sub>i</sub>* both identify a character large object type, then either *PI<sub>i</sub>* and *P2<sub>i</sub>* shall both be character large object locator parameters or neither shall be character large object locator parameters.
- c) If *DT1<sub>i</sub>* and *DT2<sub>i</sub>* both identify a user-defined type, then either *PI<sub>i</sub>* and *P2<sub>i</sub>* shall both be user-defined type locator parameters or neither shall be user-defined type locator parameters.

## Access Rules

*None.*

## General Rules

- 1) If the SQL-agent that performs a call of an <externally-invoked procedure> in an <SQL-client module definition> is not a program that conforms to the specification for the programming language specified by the <language clause> of that <SQL-client module definition>, then the effect is implementation-dependent.
- 2) If the SQL-agent performs calls of <externally-invoked procedure>s from more than one Ada task, then the results are implementation-dependent.
- 3) If FOR STATIC ONLY is specified, then the SQL-client module includes an indication that prepared statements resulting from execution of externally-invoked procedures included in that module have no owner.
- 4) After the last time that an SQL-agent performs a call of an <externally-invoked procedure>:
  - a) 09A <rollback statement> or a <commit statement> is effectively executed. If an unrecoverable error has occurred, or if the SQL-agent terminated unexpectedly, or if any constraint is not satisfied, then a <rollback statement> is performed. Otherwise, the choice of which of these SQL-statements to perform is implementation-dependent. If the implementation choice is <commit statement>, then all holdable cursors are first closed. The determination of whether an SQL-agent has terminated unexpectedly is implementation-dependent.
  - b) For every SQL descriptor area that is currently allocated within an SQL-session associated with the SQL-agent, let *D* be the <descriptor name> of that SQL descriptor area; a <deallocate descriptor statement> that specifies
 

```
DEALLOCATE DESCRIPTOR D
```

 is effectively executed.
  - c) All SQL-sessions associated with the SQL-agent are terminated.

## Conformance Rules

- 1) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <module path specification>.
- 2) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <module transform group specification>.
- 3) Without Feature F693, “SQL-session and client module collations”, conforming SQL language shall not contain a <module collation specification>.
- 4) Without Feature B051, “Enhanced execution rights”, conforming SQL language shall not contain a <module authorization clause> that immediately contains FOR STATIC ONLY or FOR STATIC AND DYNAMIC.
- 5) Without Feature B111, “Module language Ada”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains ADA.
- 6) Without Feature B112, “Module language C”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains C.
- 7) Without Feature B113, “Module language COBOL”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains COBOL.

**13.1 <SQL-client module definition>**

- 8) Without Feature B114, “Module language Fortran”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains FORTRAN.
- 9) Without Feature B115, “Module language MUMPS”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains M.
- 10) Without Feature B116, “Module language Pascal”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains PASCAL.
- 11) Without Feature B117, “Module language PL/I”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains PLI.

## 13.2 <module name clause>

### Function

Name an SQL-client module.

### Format

```
<module name clause> ::=
 MODULE [<SQL-client module name>] [<module character set specification>]

<module character set specification> ::=
 NAMES ARE <character set specification>
```

### Syntax Rules

- 1) If a <module name clause> does not specify an <SQL-client module name>, then the <SQL-client module definition> is unnamed.
- 2) The <SQL-client module name> shall not be equivalent to the <SQL-client module name> of any other <SQL-client module definition> in the same SQL-environment.  
 NOTE 364 — An SQL-environment may have multiple <SQL-client module definition>s that are unnamed.
- 3) If the <language clause> of the containing <SQL-client module definition> specifies ADA, then an <SQL-client module name> shall be specified, and that <SQL-client module name> shall be a valid Ada library unit name.
- 4) If a <module character set specification> is not specified, then a <module character set specification> that specifies an implementation-defined character set that contains at least every character that is in <SQL language character> is implicit.

### Access Rules

*None.*

### General Rules

- 1) If <SQL-client module name> is specified, then, in the SQL-environment, the containing <SQL-client module definition> has the name given by <SQL-client module name>.

### Conformance Rules

- 1) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <module character set specification>.

## 13.3 <externally-invoked procedure>

This Subclause is modified by *Subclause 14.2*, “<externally-invoked procedure>”, in ISO/IEC 9075-9.

### Function

Define an externally-invoked procedure.

### Format

```
<externally-invoked procedure> ::=
 PROCEDURE <procedure name> <host parameter declaration list> <semicolon>
 <SQL procedure statement> <semicolon>

<host parameter declaration list> ::=
 <left paren> <host parameter declaration>
 [{ <comma> <host parameter declaration> }...] <right paren>

<host parameter declaration> ::=
 <host parameter name> <host parameter data type>
 | <status parameter>

<host parameter data type> ::=
 <data type> [<locator indication>]

<status parameter> ::=
 SQLSTATE
```

### Syntax Rules

- 1) 09 The <procedure name> shall not be equivalent to the <procedure name> of any other <externally-invoked procedure> in the containing <SQL-client module definition>.
 

NOTE 365 — The <procedure name> should be a standard-conforming procedure, function, or routine name of the language specified by the subject <language clause>. Failure to observe this recommendation will have implementation-dependent effects.
- 2) The <host parameter name> of each <host parameter declaration> in an <externally-invoked procedure> shall not be equivalent to the <host parameter name> of any other <host parameter declaration> in that <externally-invoked procedure>.
- 3) Any <host parameter name> contained in the <SQL procedure statement> of an <externally-invoked procedure> shall be specified in a <host parameter declaration> in that <externally-invoked procedure>.
- 4) If <locator indication> is simply contained in <host parameter declaration>, then:
  - a) The declared type *T* identified by the <data type> immediately contained in <host parameter data type> shall be either binary large object type, character large object type, array type, multiset type, or user-defined type.
  - b) If *T* is a binary large object type, then the host parameter identified by <host parameter name> is called a *binary large object locator parameter*.

**13.3 <externally-invoked procedure>**

- c) If *T* is a character large object type, then the host parameter identified by <host parameter name> is called a *character large object locator parameter*.
  - d) If *T* is an array type, then the host parameter identified by <host parameter name> is called an *array locator parameter*.
  - e) If *T* is a multiset type, then the host parameter identified by <host parameter name> is called a *multiset locator parameter*.
  - f) If *T* is a user-defined type, then the host parameter identified by <host parameter name> is called a *user-defined type locator parameter*.
- 5) A call of an <externally-invoked procedure> shall supply *n* arguments, where *n* is the number of <host parameter declaration>s in the <externally-invoked procedure>.
  - 6) An <externally-invoked procedure> shall contain one <status parameter> referred to as an *SQLSTATE host parameter*. The SQLSTATE host parameter is referred to as a *status parameter*.
  - 7) For each <host parameter declaration>, to determine whether the corresponding host parameter is an input host parameter, an output host parameter, or both an input host parameter and an output host parameter, the Syntax Rules of Subclause 9.6, “Host parameter mode determination”, are applied with <host parameter declaration> as *HOST PARAM DECL* and <SQL procedure statement> as *SQL PROC STMT*.
  - 8) The Syntax Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, shall be satisfied.

**Access Rules**

*None.*

**General Rules**

- 1) An <externally-invoked procedure> defines an *externally-invoked procedure* that may be called by an SQL-agent.
- 2) If the <SQL-client module definition> that contains the <externally-invoked procedure> is associated with an SQL-agent that is associated with another <SQL-client module definition> that contains an <externally-invoked procedure> with equivalent <procedure name>s, then the effect is implementation-defined.
- 3) The language identified by the <language name> contained in the <language clause> of the <SQL-client module definition> that contains an <externally-invoked procedure> is the *caller language* of the <externally-invoked procedure>.
- 4) If the SQL-agent that performs a call of a <externally-invoked procedure> is not a program that conforms to the specification for the programming language specified by the caller language of the <externally-invoked procedure>, then the effect is implementation-dependent.
- 5) If the caller language of an <externally-invoked procedure> is ADA and the SQL-agent performs calls of <externally-invoked procedure>s from more than one Ada task, then the results are implementation-dependent.
- 6) If the <SQL-client module definition> that contains the <externally-invoked procedure> has an explicit <module authorization identifier> *MAI* that is not equivalent to the SQL-session <authorization identifier> *SAI*, then:

**13.3 <externally-invoked procedure>**

- a) Whether or not *SAI* can invoke <externally-invoked procedure>s in an <SQL-client module definition> with explicit <module authorization identifier> *MAI* is implementation-defined, as are any restrictions pertaining to such invocation.
  - b) If *SAI* is restricted from invoking an <externally-invoked procedure> in an <SQL-client module definition> with explicit <module authorization identifier> *MAI*, then an exception condition is raised: *invalid authorization specification*.
- 7) If the value of any input host parameter provided by the SQL-agent falls outside the set of allowed values of the declared type of the host parameter, or if the value of any output host parameter resulting from the execution of the <externally-invoked procedure> falls outside the set of values supported by the SQL-agent for that host parameter, then the effect is implementation-defined. If the implementation-defined effect is the raising of an exception condition, then an exception condition is raised: *data exception — invalid parameter value*.
  - 8) A copy of the top cell of the authorization stack is pushed onto the authorization stack. If the SQL-client module *M* that includes the externally-invoked procedure has an owner, then the top cell of the authorization stack is set to contain only the authorization identifier of the owner of *M*.
  - 9) If this is the first invocation of any externally-invoked procedure of *M* in the current SQL-session, then for each cursor declaration descriptor *CDD* of a declared cursor of *M*, let *CID* be a new cursor instance descriptor whose cursor declaration descriptor is *CDD*, whose SQL-session is the current SQL-session, and whose state is closed.
  - 10) Let *S* be the <SQL procedure statement> of the <externally-invoked procedure>.
  - 11) The General Rules of Subclause 13.5, “<SQL procedure statement>”, are evaluated with *S* as the executing statement.
  - 12) Upon completion of execution, the top cell in the authorization stack is removed.

**Conformance Rules**

- 1) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <host parameter data type> that simply contains a <data type> that specifies a structured type and that contains a <locator indication>.
- 2) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <host parameter data type> that simply contains an <array type> and that contains a <locator indication>.
- 3) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <host parameter data type> that simply contains a <multiset type> and that contains a <locator indication>.

## 13.4 Calls to an <externally-invoked procedure>

*This Subclause is modified by Subclause 11.1, “Calls to an <externally-invoked procedure>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 14.3, “Calls to an <externally-invoked procedure>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 13.1, “Calls to an <externally-invoked procedure>”, in ISO/IEC 9075-14.*

### Function

Define the call to an <externally-invoked procedure> by an SQL-agent.

### Syntax Rules

- 1) Let  $n$  be the number of <host parameter declaration>s in the <externally-invoked procedure>  $EP$  being called. Let  $PD_i$ ,  $1 \text{ (one)} \leq i \leq n$ , be the  $i$ -th <host parameter declaration>. Let  $PDT_i$  be the <data type> contained in  $PD_i$ .
- 2) If the caller language of the <externally-invoked procedure> is ADA, then:
  - a) The SQL-implementation shall generate the source code of an Ada library unit package  $ALUP$  the name of which shall be
 

Case:

    - i) If the <SQL-client module name>  $SCMN$  of the <SQL-client module definition> is a valid Ada identifier, then equivalent to  $SCMN$ .
    - ii) Otherwise, implementation-defined.
  - b) For each <externally-invoked procedure> of the <SQL-client module definition>, there shall appear within  $ALUP$  a subprogram declaration declaring a procedure.
    - i) If <procedure name> is a valid Ada identifier, then the name of that procedure  $PN$  shall be equivalent to <procedure name>; otherwise,  $PN$  shall be implementation-defined.
    - ii) The parameters in each Ada procedure declaration  $APD$  shall appear in the same order as the <host parameter declaration>s of the corresponding <externally-invoked procedure>  $EIP$ . If the names of the parameters declared in the <host parameter declaration>s of  $EIP$  are valid Ada identifiers, then the parameters in  $APD$  shall have parameter names that are equivalent to the names of the corresponding parameters declared in the <host parameter declaration>s contained in  $EIP$ ; otherwise, the parameters in  $APD$  shall parameter names that are implementation-defined
    - iii) The parameter modes and subtype marks used in the parameter specifications are constrained by the remaining paragraphs of this Subclause.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of Table 16, “Data type correspondences for Ada”, for which the corresponding row in the “Ada data type” column is 'None'.



## 13.4 Calls to an &lt;externally-invoked procedure&gt;

- d) The types of parameter specifications within the Ada subprogram declarations shall be taken from the library unit package `Interfaces.SQL` and its children `Numerics` and `Varying` and optional children `Adacsn` and `Adacsn.Varying`.
- e) 040914 The declaration of the library unit package `Interfaces.SQL` shall conform to the following template:

```

package Interfaces.SQL is
- The declarations of CHAR and NCHAR may be subtype declarations
 type CHAR is (See the Syntax Rules)
 type NCHAR is (See the Syntax Rules)
 type SMALLINT is range bs .. ts;
 type INT is range bi .. ti;
 type BIGINT is range bb .. tb;
 type REAL is digits dr;
 type DOUBLE_PRECISION is digits dd;
 type BOOLEAN is new Boolean;
 subtype INDICATOR_TYPE is t;
 type SQLSTATE_TYPE is new CHAR (1 .. 5);
package SQLSTATE_CODES is
 AMBIGUOUS_CURSOR_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "3C000";
 ATTEMPT_TO_ASSIGN_TO_NON_UPDATABLE_COLUMN_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0U000";
 ATTEMPT_TO_ASSIGN_TO_ORDERING_COLUMN_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0V000";
 CARDINALITY_VIOLATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "21000";
 CLI_SPECIFIC_CONDITION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "HY000";
 CONNECTION_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "08000";
 CONNECTION_EXCEPTION_CONNECTION_DOES_NOT_EXIST:
 constant SQLSTATE_TYPE := "08003";
 CONNECTION_EXCEPTION_CONNECTION_FAILURE:
 constant SQLSTATE_TYPE := "08006";
 CONNECTION_EXCEPTION_CONNECTION_NAME_IN_USE:
 constant SQLSTATE_TYPE := "08002";
 CONNECTION_EXCEPTION_SQLCLIENT_UNABLE_TO_ESTABLISH_SQLCONNECTION:
 constant SQLSTATE_TYPE := "08001";
 CONNECTION_EXCEPTION_SQLSERVER_REJECTED_ESTABLISHMENT_OF_SQLCONNECTION:
 constant SQLSTATE_TYPE := "08004";
 CONNECTION_EXCEPTION_TRANSACTION_RESOLUTION_UNKNOWN:
 constant SQLSTATE_TYPE := "08007";
 DATA_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "22000";
 DATA_EXCEPTION_ARRAY_ELEMENT_ERROR:
 constant SQLSTATE_TYPE := "2202E";
 DATA_EXCEPTION_ATTEMPT_TO_REPLACE_A_ZERO_LENGTH_STRING:
 constant SQLSTATE_TYPE := "2201U";
 DATA_EXCEPTION_CHARACTER_NOT_IN_REPERTOIRE:
 constant SQLSTATE_TYPE := "22021";
 DATA_EXCEPTION_DATETIME_FIELD_OVERFLOW:
 constant SQLSTATE_TYPE := "22008";
 DATA_EXCEPTION_DIVISION_BY_ZERO:
 constant SQLSTATE_TYPE := "22012";
 DATA_EXCEPTION_ERROR_IN_ASSIGNMENT:

```

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

```

 constant SQLSTATE_TYPE := "22005";
DATA_EXCEPTION_ESCAPE_CHARACTER_CONFLICT:
 constant SQLSTATE_TYPE := "2200B";
DATA_EXCEPTION_INDICATOR_OVERFLOW:
 constant SQLSTATE_TYPE := "22022";
DATA_EXCEPTION_INTERVAL_FIELD_OVERFLOW:
 constant SQLSTATE_TYPE := "22015";
DATA_EXCEPTION_INTERVAL_VALUE_OUT_OF_RANGE:
 constant SQLSTATE_TYPE := "2200P";
DATA_EXCEPTION_INVALID_ARGUMENT_FOR_NATURAL_LOGARITHM:
 constant SQLSTATE_TYPE := "2201E";
DATA_EXCEPTION_INVALID_ARGUMENT_FOR_POWER_FUNCTION:
 constant SQLSTATE_TYPE := "2201F";
DATA_EXCEPTION_INVALID_ARGUMENT_FOR_WIDTH_BUCKET_FUNCTION:
 constant SQLSTATE_TYPE := "2201G";
DATA_EXCEPTION_INVALID_CHARACTER_VALUE_FOR_CAST:
 constant SQLSTATE_TYPE := "22018";
DATA_EXCEPTION_INVALID_DATETIME_FORMAT:
 constant SQLSTATE_TYPE := "22007";
DATA_EXCEPTION_INVALID_ESCAPE_CHARACTER:
 constant SQLSTATE_TYPE := "22019";
DATA_EXCEPTION_INVALID_ESCAPE_OCTET:
 constant SQLSTATE_TYPE := "2200D";
DATA_EXCEPTION_INVALID_ESCAPE_SEQUENCE:
 constant SQLSTATE_TYPE := "22025";
DATA_EXCEPTION_INVALID_INDICATOR_PARAMETER_VALUE:
 constant SQLSTATE_TYPE := "22010";
DATA_EXCEPTION_INVALID_INTERVAL_FORMAT:
 constant SQLSTATE_TYPE := "22006";
DATA_EXCEPTION_INVALID_PARAMETER_VALUE:
 constant SQLSTATE_TYPE := "22023";
DATA_EXCEPTION_INVALID_PRECEDING_OR_FOLLOWING_SIZE_IN_WINDOW_FUNCTION:
 constant SQLSTATE_TYPE := "22013";
DATA_EXCEPTION_INVALID_REGULAR_EXPRESSION:
 constant SQLSTATE_TYPE := "2201B";
DATA_EXCEPTION_INVALID_REPEAT_ARGUMENT_IN_A_SAMPLE_CLAUSE:
 constant SQLSTATE_TYPE := "2202G";
DATA_EXCEPTION_INVALID_SAMPLE_SIZE:
 constant SQLSTATE_TYPE := "2202H";
DATA_EXCEPTION_INVALID_TIME_ZONE_DISPLACEMENT_VALUE:
 constant SQLSTATE_TYPE := "22009";
DATA_EXCEPTION_INVALID_USE_OF_ESCAPE_CHARACTER:
 constant SQLSTATE_TYPE := "2200C";
DATA_EXCEPTION_INVALID_XQUERY_OPTION_FLAG:
 constant SQLSTATE_TYPE := "2201T";
DATA_EXCEPTION_INVALID_XQUERY_REGULAR_EXPRESSION:
 constant SQLSTATE_TYPE := "2201S";
DATA_EXCEPTION_INVALID_XQUERY_REPLACEMENT_STRING:
 constant SQLSTATE_TYPE := "2201V";
DATA_EXCEPTION_NULL_VALUE_NO_INDICATOR_PARAMETER:
 constant SQLSTATE_TYPE := "2200G";
DATA_EXCEPTION_MOST_SPECIFIC_TYPE_MISMATCH:
 constant SQLSTATE_TYPE := "22002";
DATA_EXCEPTION_MULTISSET_VALUE_OVERFLOW:
 constant SQLSTATE_TYPE := "2200Q";
DATA_EXCEPTION_NONCHARACTER_IN_UCS_STRING:
 constant SQLSTATE_TYPE := "22029";

```

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

```

DATA_EXCEPTION_NULL_VALUE_NOT_ALLOWED:
 constant SQLSTATE_TYPE := "22004";
DATA_EXCEPTION_NULL_VALUE_SUBSTITUTED_FOR_MUTATOR_SUBJECT_PARAMETER:
 constant SQLSTATE_TYPE := "2202D";
DATA_EXCEPTION_NUMERIC_VALUE_OUT_OF_RANGE:
 constant SQLSTATE_TYPE := "22003";
DATA_EXCEPTION_SEQUENCE_GENERATOR_LIMIT_EXCEEDED:
 constant SQLSTATE_TYPE := "2200H";
DATA_EXCEPTION_STRING_DATA_LENGTH_MISMATCH:
 constant SQLSTATE_TYPE := "22026";
DATA_EXCEPTION_STRING_DATA_RIGHT_TRUNCATION:
 constant SQLSTATE_TYPE := "22001";
DATA_EXCEPTION_SUBSTRING_ERROR:
 constant SQLSTATE_TYPE := "22011";
DATA_EXCEPTION_TRIM_ERROR:
 constant SQLSTATE_TYPE := "22027";
DATA_EXCEPTION_UNTERMINATED_C_STRING:
 constant SQLSTATE_TYPE := "22024";
DATA_EXCEPTION_ZERO_LENGTH_CHARACTER_STRING:
 constant SQLSTATE_TYPE := "2200F";
DEPENDENT_PRIVILEGE_DESCRIPTOR_STILL_EXIST_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2B000";
DIAGNOSTICS_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0Z000";
DIAGNOSTICS_EXCEPTION_MAXIMUM_NUMBER_OF_DIAGNOSTICS_AREAS_EXCEEDED:
 constant SQLSTATE_TYPE := "0Z001";
DYNAMIC_SQL_ERROR_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "07000";
DYNAMIC_SQL_ERROR_CURSOR_SPECIFICATION_CANNOT_BE_EXECUTED:
 constant SQLSTATE_TYPE := "07003";
DYNAMIC_SQL_ERROR_INVALID_DATETIME_INTERVAL_CODE:
 constant SQLSTATE_TYPE := "0700F";
DYNAMIC_SQL_ERROR_INVALID_DESCRIPTOR_COUNT:
 constant SQLSTATE_TYPE := "07008";
DYNAMIC_SQL_ERROR_INVALID_DESCRIPTOR_INDEX:
 constant SQLSTATE_TYPE := "07009";
DYNAMIC_SQL_ERROR_PREPARED_STATEMENT_NOT_A_CURSOR_SPECIFICATION:
 constant SQLSTATE_TYPE := "07005";
DYNAMIC_SQL_ERROR_RESTRICTED_DATA_TYPE_ATTRIBUTE_VIOLATION:
 constant SQLSTATE_TYPE := "07006";
DYNAMIC_SQL_ERROR_DATA_TYPE_TRANSFORM_FUNCTION_VIOLATION:
 constant SQLSTATE_TYPE := "0700B";
DYNAMIC_SQL_ERROR_INVALID_DATA_TARGET:
 constant SQLSTATE_TYPE := "0700D";
DYNAMIC_SQL_ERROR_INVALID_LEVEL_VALUE:
 constant SQLSTATE_TYPE := "0700E";
DYNAMIC_SQL_ERROR_UNDEFINED_DATA_VALUE:
 constant SQLSTATE_TYPE := "0700C";
DYNAMIC_SQL_ERROR_USING_CLAUSE_DOES_NOT_MATCH_DYNAMIC_PARAMETER_SPEC:
 constant SQLSTATE_TYPE := "07001";
DYNAMIC_SQL_ERROR_USING_CLAUSE_DOES_NOT_MATCH_TARGET_SPEC:
 constant SQLSTATE_TYPE := "07002";
DYNAMIC_SQL_ERROR_USING_CLAUSE_REQUIRED_FOR_DYNAMIC_PARAMETERS:
 constant SQLSTATE_TYPE := "07004";
DYNAMIC_SQL_ERROR_USING_CLAUSE_REQUIRED_FOR_RESULT_FIELDS:
 constant SQLSTATE_TYPE := "07007";
EXTERNAL_ROUTINE_EXCEPTION_NO_SUBCLASS:

```

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

```

 constant SQLSTATE_TYPE := "38000";
EXTERNAL_ROUTINE_EXCEPTION_CONTAINING_SQL_NOT_PERMITTED:
 constant SQLSTATE_TYPE := "38001";
EXTERNAL_ROUTINE_EXCEPTION_MODIFYING_SQL_DATA_NOT_PERMITTED:
 constant SQLSTATE_TYPE := "38002";
EXTERNAL_ROUTINE_EXCEPTION_PROHIBITED_SQL_STATEMENT_ATTEMPTED:
 constant SQLSTATE_TYPE := "38003";
EXTERNAL_ROUTINE_EXCEPTION_READING_SQL_DATA_NOT_PERMITTED:
 constant SQLSTATE_TYPE := "38004";
EXTERNAL_ROUTINE_INVOCATION_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "39000";
EXTERNAL_ROUTINE_INVOCATION_EXCEPTION_NULL_VALUE_NOT_ALLOWED:
 constant SQLSTATE_TYPE := "39004";
FEATURE_NOT_SUPPORTED_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0A000";
FEATURE_NOT_SUPPORTED_MULTIPLE_ENVIRONMENT_TRANSACTIONS:
 constant SQLSTATE_TYPE := "0A001";
INTEGRITY_CONSTRAINT_VIOLATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "23000";
INTEGRITY_CONSTRAINT_VIOLATION_RESTRICT_VIOLATION:
 constant SQLSTATE_TYPE := "23001";
INVALID_AUTHORIZATION_SPECIFICATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "28000";
INVALID_CATALOG_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "3D000";
INVALID_CHARACTER_SET_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2C000";
INVALID_COLLATION_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2H000";
INVALID_CONDITION_NUMBER_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "35000";
INVALID_CONNECTION_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2E000";
INVALID_CURSOR_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "34000";
INVALID_CURSOR_STATE_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "24000";
INVALID_GRANTOR_STATE_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0L000";
INVALID_ROLE_SPECIFICATION:
 constant SQLSTATE_TYPE := "0P000";
INVALID_SCHEMA_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "3F000";
INVALID_SCHEMA_NAME_LIST_SPECIFICATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0E000";
INVALID_SQL_DESCRIPTOR_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "33000";
INVALID_SQL_INVOKED_PROCEDURE_REFERENCE_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0M000";
INVALID_SQL_STATEMENT:
 constant SQLSTATE_TYPE := "30000";
INVALID_SQL_STATEMENT_IDENTIFIER_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "30000";
INVALID_SQL_STATEMENT_NAME_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "26000";
INVALID_TRANSFORM_GROUP_NAME_SPECIFICATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0S000";

```

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

```

INVALID_TRANSACTION_STATE_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "25000";
INVALID_TRANSACTION_STATE_ACTIVE_SQL_TRANSACTION:
 constant SQLSTATE_TYPE := "25001";
INVALID_TRANSACTION_STATE_BRANCH_TRANSACTION_ALREADY_ACTIVE:
 constant SQLSTATE_TYPE := "25002";
INVALID_TRANSACTION_STATE_HELD_CURSOR_REQUIRES_SAME_ISOLATION_LEVEL:
 constant SQLSTATE_TYPE := "25008";
INVALID_TRANSACTION_STATE_INAPPROPRIATE_ACCESS_MODE_FOR_BRANCH_TRANSACTION:
 constant SQLSTATE_TYPE := "25003";
INVALID_TRANSACTION_STATE_INAPPROPRIATE_ISOLATION_LEVEL_FOR_BRANCH_TRANSACTION:

 constant SQLSTATE_TYPE := "25004";
INVALID_TRANSACTION_STATE_NO_ACTIVE_SQL_TRANSACTION_FOR_BRANCH_TRANSACTION:
 constant SQLSTATE_TYPE := "25005";
INVALID_TRANSACTION_STATE_READ_ONLY_SQL_TRANSACTION:
 constant SQLSTATE_TYPE := "25006";
INVALID_TRANSACTION_STATE_SCHEMA_AND_DATA_STATEMENT_MIXING_NOT_SUPPORTED:
 constant SQLSTATE_TYPE := "25007";
INVALID_TRANSACTION_TERMINATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2D000";
LOCATOR_EXCEPTION_INVALID_SPECIFICATION:
 constant SQLSTATE_TYPE := "0F001";
LOCATOR_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0F000";
NO_DATA_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "02000";
NO_DATA_NO_ADDITIONAL_DYNAMIC_RESULT_SETS_RETURNED:
 constant SQLSTATE_TYPE := "02001";
NO_DATA_NO_DYNAMIC_RESULT_SETS_RETURNED:
 constant SQLSTATE_TYPE := "02001";
REMOTE_DATABASE_ACCESS_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "HZ000";
SAVEPOINT_EXCEPTION_INVALID_SPECIFICATION:
 constant SQLSTATE_TYPE := "3B001";
SAVEPOINT_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "3B000";
SAVEPOINT_EXCEPTION_TOO_MANY:
 constant SQLSTATE_TYPE := "3B002";
SQL_ROUTINE_EXCEPTION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2F000";
SQL_ROUTINE_EXCEPTION_FUNCTION_EXECUTED_NO_RETURN_STATEMENT:
 constant SQLSTATE_TYPE := "2F005";
SQL_ROUTINE_EXCEPTION_MODIFYING_SQL_DATA_NOT_PERMITTED:
 constant SQLSTATE_TYPE := "2F002";
SQL_ROUTINE_EXCEPTION_PROHIBITED_SQL_STATEMENT_ATTEMPTED:
 constant SQLSTATE_TYPE := "2F003";
SQL_ROUTINE_EXCEPTION_READING_SQL_DATA_NOT_PERMITTED:
 constant SQLSTATE_TYPE := "2F004";
SUCCESSFUL_COMPLETION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "00000";
SYNTAX_ERROR_OR_ACCESS_RULE_VIOLATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "42000";
SYNTAX_ERROR_OR_ACCESS_RULE_VIOLATION_IN_DIRECT_STATEMENT_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "2A000";
SYNTAX_ERROR_OR_ACCESS_RULE_VIOLATION_IN_DYNAMIC_STATEMENT_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "37000";

```

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

```

TARGET_TABLE_DISAGREES_WITH_CURSOR_SPECIFICATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "0T000";
TRANSACTION_ROLLBACK_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "40000";
TRANSACTION_ROLLBACK_INTEGRITY_CONSTRAINT_VIOLATION:
 constant SQLSTATE_TYPE := "40002";
TRANSACTION_ROLLBACK_SERIALIZATION_FAILURE:
 constant SQLSTATE_TYPE := "40001";
TRANSACTION_ROLLBACK_STATEMENT_COMPLETION_UNKNOWN:
 constant SQLSTATE_TYPE := "40003";
TRIGGERED_DATA_CHANGE_VIOLATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "27000";
WARNING_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "01000";
WARNING_ADDITIONAL_RESULT_SETS_RETURNED:
 constant SQLSTATE_TYPE := "0100D";
WARNING_ARRAY_DATA_RIGHT_TRUNCATION:
 constant SQLSTATE_TYPE := "0102F";
WARNING_ATTEMPT_TO_RETURN_TOO_MANY_RESULT_SETS:
 constant SQLSTATE_TYPE := "0100E";
WARNING_CURSOR_OPERATION_CONFLICT:
 constant SQLSTATE_TYPE := "01001";
WARNING_DEFAULT_VALUE_TOO_LONG_FOR_INFORMATION_SCHEMA:
 constant SQLSTATE_TYPE := "0100B";
WARNING_DISCONNECT_ERROR:
 constant SQLSTATE_TYPE := "01002";
WARNING_DYNAMIC_RESULT_SETS_RETURNED:
 constant SQLSTATE_TYPE := "0100C";
WARNING_INSUFFICIENT_ITEM_DESCRIPTOR_AREAS:
 constant SQLSTATE_TYPE := "01005";
WARNING_NULL_VALUE_ELIMINATED_IN_SET_FUNCTION:
 constant SQLSTATE_TYPE := "01003";
WARNING_PRIVILEGE_NOT_GRANTED:
 constant SQLSTATE_TYPE := "01007";
WARNING_PRIVILEGE_NOT_REVOKED:
 constant SQLSTATE_TYPE := "01006";
WARNING_QUERY_EXPRESSION_TOO_LONG_FOR_INFORMATION_SCHEMA:
 constant SQLSTATE_TYPE := "0100A";
WARNING_RESULT_SETS_RETURNED:
 constant SQLSTATE_TYPE := "0100C";
WARNING_SEARCH_CONDITION_TOO_LONG_FOR_INFORMATION_SCHEMA:
 constant SQLSTATE_TYPE := "01009";
WARNING_STATEMENT_TOO_LONG_FOR_INFORMATION_SCHEMA:
 constant SQLSTATE_TYPE := "0100F";
WARNING_STRING_DATA_RIGHT_TRUNCATION_WARNING:
 constant SQLSTATE_TYPE := "01004";
WITH_CHECK_OPTION_VIOLATION_NO_SUBCLASS:
 constant SQLSTATE_TYPE := "44000";
end SQLSTATE_CODES;
end Interfaces.SQL;

```

where *bs*, *ts*, *bi*, *ti*, *bb*, *tb*, *dr*, *dd*, *bsc*, and *tsc* are implementation-defined integer values. *t* is INT or SMALLINT, corresponding with an implementation-defined <exact numeric type> of indicator parameters.

NOTE 366 — The Ada identifier INVALID\_SQL\_STATEMENT appears for compatibility with earlier editions of ISO/IEC 9075. However, the intended symbol is INVALID\_SQL\_STATEMENT\_IDENTIFIER\_NO\_SUBCLASS, which was added to correspond correctly with the exception condition name.

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

- f) The library unit package `Interfaces.SQL.Numerics` shall contain a sequence of decimal fixed point type declarations of the following form.

```
type Scale_s is delta 10.0 ** - s digits max_p;
```

where *s* is an integer ranging from 0 (zero) to an implementation-defined maximum value and *max\_p* is an implementation-defined integer maximum precision.

- g) The library unit package `Interfaces.SQL.Varying` shall contain type or subtype declarations with the defining identifiers `CHAR` and `NCHAR`.
- h) Let *SQLcsn* be a <character set name> and let *Adacsn* be the result of replacing <period>s in *SQLcsn* with <underscore>s. If *Adacsn* is a valid Ada identifier, then the library unit packages `Interfaces.SQL.Adacsn` and `Interfaces.SQL.Adacsn.Varying` shall contain a type or subtype declaration with defining identifier `CHAR`. If *Adacsn* is not a valid Ada identifier, then the names of these packages shall be implementation-defined.
- i) `Interfaces.SQL` and its children may contain context clauses and representation items as needed. These packages may also contain declarations of Ada character types as needed to support the declarations of the types `CHAR` and `NCHAR`.

NOTE 367 — If the implementation-defined character set specification used by default with a fixed-length character string type is Latin1, then the declaration

```
subtype CHAR is String;
```

within `Interfaces.SQL` and the declaration

```
subtype CHAR is
 Ada.Strings.Unbounded.Unbounded_String;
```

within `Interfaces.SQL.Varying` (assuming the appropriate context clause) conform to the requirements of this paragraph of this Subclause. If the character set underlying NATIONAL CHARACTER is supported by an Ada package specification `Host_Char_Pkg` that declares a type `String_Type` that stores strings over the given character set, and furthermore the package specification `Host_Char_Pkg_Varying` (not necessary distinct from `Host_Char_Pkg`) declares a type `String_Type_Varying` that reproduces the functionality of `Ada.Strings.Unbounded.Unbounded_String` over the national character string type (rather than Latin1), then the declaration

```
subtype NCHAR is Host_Char_Pkg.String_Type;
```

within `Interfaces.SQL` and the declaration

```
subtype NCHAR is Host_Char_Pkg_Varying.String_Type_Varying;
```

within `Interfaces.SQL.Varying` conform to the requirements of this paragraph. Similar comments apply to other character sets and the packages `Interfaces.SQL.Adacsn` and `Interfaces.SQL.Adacsn.Varying`.

- j) The library unit package `Interfaces.SQL` shall contain declarations of the following form:

```
package CHARACTER_SET renames Interfaces.SQL.Adacsn;
subtype CHARACTER_TYPE is CHARACTER_SET.cst;
```

where *cst* is a data type capable of storing a single character from the default character set. The package `Interfaces.SQL.Adacsn` shall contain the necessary declaration for *cst*.

NOTE 368 — If the default character set is Latin1, then a declaration of the form:

```
package CHARACTER_SET is
```

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

```

 subtype cst is Character;
end CHARACTER_SET;

```

may be substituted for the renaming declaration of CHARACTER\_SET.

- k) The base type of the SQLSTATE parameter shall be `Interfaces.SQL.SQLSTATE_TYPE`.
- l) The Ada parameter mode of the SQLSTATE parameter is out.
- m) If the *i*-th <host parameter declaration> specifies a <data type> that is:
  - i) CHARACTER(*L*) for some *L*, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.CHAR`.
  - ii) CHARACTER VARYING(*L*) for some *L*, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.VARYING.CHAR`.
  - iii) NATIONAL CHARACTER(*L*) for some *L*, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.NCHAR`.
  - iv) NATIONAL CHARACTER VARYING(*L*) for some *L*, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.VARYING.NCHAR`.
  - v) CHARACTER(*L*) CHARACTER SET *csn* for some *L* and some character set name *csn*, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.Adacsn.CHAR`.
  - vi) CHARACTER VARYING(*L*) CHARACTER SET *csn* for some *L* and some character set name *csn*, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.Adacsn.VARYING.CHAR`.

If *P* is an actual parameter associated with the *i*-th parameter in a call to the encompassing procedure, then *P* shall be sufficient to hold a character string of length *L* in the appropriate character set.

NOTE 369 — If a character set uses fixed length encodings then the definition of the subtype CHAR for fixed length strings may be an array type whose element type is an Ada character type. If that Ada character type is defined so as to use the number of bits per character used by the SQL encoding, then the restriction on *P* is precisely `P'LENGTH = L`. For variable length strings using fixed length encodings, if the definition of CHAR in the appropriate VARYING package is based on the type `Ada.Strings.Unbounded.Unbounded_String`, there is no restriction on *P*. Otherwise, a precise statement of the restriction on *P* is implementation-defined.

- n) If the *i*-th <host parameter declaration> specifies a <data type> that is NUMERIC(*P*,*S*) for some <precision> *P* and <scale> *S*, then the Ada library unit package generated for the encompassing module shall contain a declaration equivalent to:

```

subtype Numeric_p_s is
 Interfaces.SQL.Numerics.Scale_s digits p;

```

The subtype mark in the *i*-th parameter declaration shall specify this subtype.

- o) If the *i*-th <host parameter declaration> specifies a <data type> that is SMALLINT, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.SMALLINT`.
- p) If the *i*-th <host parameter declaration> specifies a <data type> that is INTEGER, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.INT`.
- q) If the *i*-th <host parameter declaration> specifies a <data type> that is BIGINT, then the subtype mark in the *i*-th parameter declaration shall specify `Interfaces.SQL.BIGINT`.



## 13.4 Calls to an &lt;externally-invoked procedure&gt;

- r) If the  $i$ -th <host parameter declaration> specifies a <data type> that is REAL, then the subtype mark in the  $i$ -th parameter declaration shall specify `Interfaces.SQL.REAL`.
- s) If the  $i$ -th <host parameter declaration> specifies a <data type> that is DOUBLE\_PRECISION, then the subtype mark in the  $i$ -th parameter declaration shall specify `Interfaces.SQL.DOUBLE_PRECISION`.
- t) For every parameter,

Case:

- i) If the parameter is an input parameter but not an output parameter, then the Ada parameter mode is **in**.
  - ii) If the parameter is an output parameter but not an input parameter, then the Ada parameter mode is **out**.
  - iii) If the parameter is both an input parameter and an output parameter, then the Ada parameter mode is **in out**.
  - iv) Otherwise, the Ada parameter mode is **in**, **out**, or **in out**.
- u) The following Ada library unit renaming declaration exists:

```
with Interfaces.SQL;
package SQL_Standard renames Interfaces.SQL.
```

- 3) If the caller language of the <externally-invoked procedure> is C, then:
  - a) The declared type of an SQLSTATE host parameter shall be C char with length 6.
  - b) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of Table 17, “Data type correspondences for C”, for which the corresponding row in the “C data type” column is 'None'.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the type of the  $i$ -th host parameter shall be the data type listed in the “C data type” column of Table 17, “Data type correspondences for C”, for which the corresponding row in the “SQL data type” column is  $PDT_i$ .
- 4) If the caller language of the <externally-invoked procedure> is COBOL, then:
  - a) The declared type of an SQLSTATE host parameter shall be COBOL PICTURE X(5).
  - b) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of Table 18, “Data type correspondences for COBOL”, for which the corresponding row in the “COBOL data type” column is 'None'.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the type of the  $i$ -th host parameter shall be the data type listed in the “COBOL data type” column of Table 18, “Data type correspondences for COBOL”, for which the corresponding row in the “SQL data type” column is  $PDT_i$ .
- 5) If the caller language of the <externally-invoked procedure> is FORTRAN, then:
  - a) The declared type of an SQLSTATE host parameter shall be Fortran CHARACTER with length 5.

**13.4 Calls to an <externally-invoked procedure>**

- b) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of [Table 19, “Data type correspondences for Fortran”](#), for which the corresponding row in the “Fortran data type” column is 'None'.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the type of the  $i$ -th host parameter shall be the data type listed in the “Fortran data type” column of [Table 19, “Data type correspondences for Fortran”](#), for which the corresponding row in the “SQL data type” column is  $PDT_i$ .
- 6) If the caller language of the <externally-invoked procedure> is M, then:
  - a) The declared type of an SQLSTATE host parameter shall be M character with maximum length greater than or equal to 5.
  - b) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of [Table 20, “Data type correspondences for M”](#), for which the corresponding row in the “M data type” column is 'None'.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the type of the  $i$ -th host parameter shall be the data type listed in the “M data type” column of [Table 20, “Data type correspondences for M”](#), for which the corresponding row in the “SQL data type” column is  $PDT_i$ .
- 7) If the caller language of the <externally-invoked procedure> is PASCAL, then:
  - a) The declared type of an SQLSTATE host parameter shall be Pascal PACKED ARRAY[1..5] OF CHAR.
  - b) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of [Table 21, “Data type correspondences for Pascal”](#), for which the corresponding row in the “Pascal data type” column is 'None'.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the type of the  $i$ -th host parameter shall be the data type listed in the “Pascal data type” column of [Table 21, “Data type correspondences for Pascal”](#), for which the corresponding row in the “SQL data type” column is  $PDT_i$ .
- 8) If the caller language of the <externally-invoked procedure> is PLI, then:
  - a) The declared type of an SQLSTATE host parameter shall be PL/I CHARACTER(5).
  - b) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ ,  $PDT_i$  shall not identify a data type listed in the “SQL data type” column of [Table 22, “Data type correspondences for PL/I”](#), for which the corresponding row in the “PL/I data type” column is 'None'.
  - c) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the type of the  $i$ -th host parameter shall be the data type listed in the “PL/I data type” column of [Table 22, “Data type correspondences for PL/I”](#), for which the corresponding row in the “SQL data type” column is  $PDT_i$ .

**Access Rules**

*None.*

## General Rules

- 1) Let *EP*, *PD*, *PN*, *DT*, and *PI* be a *PROC*, a *DECL*, a *NAME*, a *TYPE*, and an *ARG* specified in an application of the General Rules of this Subclause. Let *P* be the host parameter corresponding to *PD*.
- 2) If the General Rules of this Subclause are being applied for the evaluation of input parameters, and *P* is either an input host parameter or both an input host parameter and an output host parameter, then

Case:

- a) If *DT* identifies a *CHARACTER(L)* or *CHARACTER VARYING(L)* data type and the caller language of *EP* is *C*, then a reference to *PN* is implicitly treated as a character string type value in the specified character set in which the octets of *PI* are the corresponding octets of that value.

When such a reference is evaluated,

Case:

- i) If *DT* identifies a *CHARACTER(L)* data type and some *C* character preceding the least significant *C* character of the value *PI* contains the implementation-defined null character that terminates a *C* character string, then the remaining characters of the value are set to <space>s.
  - ii) If *DT* identifies a *CHARACTER VARYING(L)* data type, then the length in characters of the value is set to the number of characters of *PI* that precede the implementation-defined null character that terminates a *C* character string.
  - iii) If the least significant *C* character of the value *PI* does not contain the implementation-defined null character that terminates a *C* character string, then an exception condition is raised: *data exception — unterminated C string*; otherwise, that least significant *C* character does not correspond to any character in *PI* and is ignored.
- NOTE 370 — In each alternative of the Case, the phrase “implementation-defined null character that terminates a *C* character string” implies one or more octets all of whose bits are zero and whose number is equal to the number of octets in the largest character of the character set of *DT*.
- b) If *DT* identifies a *CHARACTER(L)* data type and the caller language of *EP* is either *COBOL*, *FORTRAN*, or *PASCAL*, or *DT* identifies a *CHARACTER VARYING(L)* data type and the caller language of *EP* is *M*, or *DT* identifies a *CHARACTER(L)* data type or *CHARACTER VARYING(L)* data type and the caller language of *EP* is *PLI*, then a reference to *PN* is implicitly treated as a character string type value in the specified character set in which the octets of *PI* are the corresponding octets of that value.
  - c) If *DT* identifies *INT*, *DEC*, or *REAL* and the caller language of *EP* is *M*, then a reference to *PN* is implicitly treated as:

CAST ( *PI* AS *DT* )

- d) If *DT* identifies a *BOOLEAN* type, then

Case:

- i) If the caller language of *EP* is *ADA*, then if *PI* is *False*, then a reference to *PN* has the value *False*; otherwise, a reference to *PN* has the value *True*.
- ii) If the caller language of *EP* is *C*, then if *PI* is 0 (zero), then a reference to *PN* has the value *False*; otherwise, a reference to *PN* has the value *True*.

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

- iii) If the caller language of *EP* is COBOL, then if *PI* is 'F', then a reference to *PN* has the value False; otherwise, a reference to *PN* has the value True.
- iv) If the caller language of *EP* is FORTRAN, then if *PI* is .FALSE., then a reference to *PN* has the value False; otherwise, a reference to *PN* has the value True.
- v) If the caller language of *EP* is PLI, then if *PI* is '0'B, then a reference to *PN* has the value False; otherwise, a reference to *PN* has the value True.

NOTE 371 — Pascal has a Boolean-type whose values are True and False.

- e) If *P* is a binary large object locator parameter, a character large object locator parameter, an array locator parameter, a multiset locator parameter, or a user-defined type locator parameter, then

Case:

- i) If *PI* is an invalid locator, then an exception condition is raised: *locator exception — invalid specification* and no further General Rules of this Subclause are applied.
  - ii) Otherwise, a reference to *PN* in a <general value specification> has the binary large object string value, the large object character string value, the array value, the multiset value, or the user-defined type value, respectively, corresponding to *PI*.
- f) If *DT* identifies a CHARACTER LARGE OBJECT or BINARY LARGE OBJECT type, then

Case:

- i) If the caller language of *EP* is C, then a reference to *PN* is implicitly treated as
- Case:
- 1) If *DT* identifies a CHARACTER LARGE OBJECT type, then a character string containing the *PN.PN\_length* characters of *PN.PN\_data* starting at character number 1 (one) in the same order that the characters appear in *PN.PN\_data*.
  - 2) If *DT* identifies a BINARY LARGE OBJECT type, then a binary large object string containing the *PN.PN\_length* octets of *PN.PN\_data* starting at octet number 1 (one) in the same order that the octets appear in *PN.PN\_data*.
- ii) If the caller language of *EP* is COBOL, then a reference to *PN* is implicitly treated as

Case:

- 1) If *DT* identifies a CHARACTER LARGE OBJECT type, then a character string containing the *PN-LENGTH OF PN* characters of *PN-DATA OF PN* starting at character number 1 (one) in the same order that the characters appear in *PN-DATA OF PN*.
  - 2) If *DT* identifies a BINARY LARGE OBJECT type, then a binary large object string containing the *PN-LENGTH OF PN* octets of *PN-DATA OF PN* starting at octet number 1 (one) in the same order that the octets appear in *PN-DATA OF PN*.
- iii) If the caller language of *EP* is FORTRAN, then a reference to *PN* is implicitly treated as

Case:

- 1) If *DT* identifies a CHARACTER LARGE OBJECT type, then a character string containing the *PN\_LENGTH* characters of *PN\_DATA* starting at character number 1 (one) in the same order that the characters appear in *PN\_DATA*.

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

- 2) If *DT* identifies a BINARY LARGE OBJECT type, then a binary large object string containing the *PN\_LENGTH* octets of *PN\_DATA* starting at octet number 1 (one) in the same order that the octets appear in *PN\_DATA*.
- iv) If the caller language of *EP* is PLI, then a reference to *PN* is implicitly treated as  
Case:
  - 1) If *DT* identifies a CHARACTER LARGE OBJECT type, then a character string containing the *PN.PN\_length* characters of *PN.PN\_data* starting at character number 1 (one) in the same order that the characters appear in *PN.PN\_data*.
  - 2) If *DT* identifies a BINARY LARGE OBJECT type, then a binary large object string containing the *PN.PN\_length* octets of *PN.PN\_data* starting at octet number 1 (one) in the same order that the octets appear in *PN.PN\_data*.
- g) If *DT* identifies a BINARY(*L*) or BINARY VARYING(*L*) data type and the caller language of *EP* is C, then a reference to *PN* is implicitly treated as a binary string type value in which the octets of *PI* are the corresponding octets of that value.

When such a reference is evaluated,

Case:

- i) If *DT* identifies a BINARY(*L*) data type and the length in octets *LPI* of *PI* is less than *L*, then a reference to the remaining *L-LPI* octets of *PN* has the value X'00'.
  - ii) If *DT* identifies a BINARY VARYING(*L*) data type, then a reference to *PN* is treated as a binary string containing the *PN.PN\_length* octets of *PN.PN\_data* starting at octet number 1 (one) in the same order that the octets appear in *PN.PN\_data*.
  - h) If *DT* identifies a BINARY(*L*) data type and the caller language of *EP* is either COBOL or FORTRAN, or *DT* identifies a BINARY(*L*) data type or BINARY VARYING(*L*) data type and the caller language of *EP* is PLI, then a reference to *PN* is implicitly treated as a binary string type value in which the octets of *PI* are the corresponding octets of that value.
  - i) Otherwise, a reference to *PN* in a <general value specification> has the value *PI*.
- 3) If the General Rules of this Subclause are being applied for the evaluation of output parameters, and *P* is either an output host parameter or both an input host parameter and an output host parameter, then

Case:

- a) If *DT* identifies CHARACTER(*L*) or CHARACTER VARYING(*L*) data types and the caller language of *EP* is C, then let *CL* be *k* greater than the maximum possible length in octets of *PN*, where *k* is the size in octets of the largest character in the character set of *DT*. A reference to *PN* that assigns some value *SV* to *PN* implicitly assigns a value that is an SQL CHARACTER(*CL*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with <space>s as necessary to reach the length *CL - k*, concatenated with a single implementation-defined null character that terminates a C character string.

NOTE 372 — The phrase “implementation-defined null character that terminates a C character string” implies one or more octets all of whose bits are zero and whose number is equal to the number of octets in the largest character of the character set of *DT*.

- b) If *DT* identifies a CHARACTER(*L*) data type and the caller language of *EP* is either COBOL, FORTRAN, or PASCAL, then let *CL* be the maximum possible length in octets of *PN*. A reference to *PN*

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

that assigns some value *SV* to *PN* implicitly assigns a value that is an SQL CHARACTER(*CL*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with <space>s as necessary to reach the length *CL*.

- c) If *DT* identifies a CHARACTER VARYING(*L*) data type and the caller language of *EP* is *M*, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns a value that is an SQL CHARACTER VARYING(*ML*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with <space>s as necessary to reach the length *CL*. *ML* is the implementation-defined maximum length of variable-length character strings.
- d) If *DT* identifies a CHARACTER(*L*) or CHARACTER VARYING(*L*) data types and the caller language of *EP* is *PLI*, then let *CL* be the maximum possible length in octets of *PN*. A reference to *PN* that assigns some value *SV* to *PN* implicitly assigns a value that is

Case:

- i) If *DT* identifies CHARACTER(*L*), then an SQL CHARACTER(*CL*) data type in which octets of the value are the corresponding octets of *SV*.
  - ii) Otherwise, an SQL CHARACTER VARYING(*CL*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with <space>s as necessary to reach the length *CL*.
- e) If *DT* identifies INT, DEC, or REAL and the caller language of *EP* is *M*, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value

CAST ( *SV* AS CHARACTER VARYING(*ML*) )

to *PI*, where *ML* is the implementation-defined maximum length of variable-length character strings.

- f) If *DT* identifies a BOOLEAN type, then

Case:

- i) If the caller language of *EP* is ADA, then a reference to *PN* that assigns the value False to *PN* implicitly assigns the value False to *PI*; a reference to *PN* that assigns the value True implicitly assigns the value True to *PI*.
- ii) If the caller language of *EP* is C, then a reference to *PN* that assigns the value False to *PN* implicitly assigns the value 0 (zero) to *PI*; a reference to *PN* that assigns the value True implicitly assigns the value 1 (one) to *PI*.
- iii) If the caller language of *EP* is COBOL, then a reference to *PN* that assigns the value False to *PN* implicitly assigns the value 'F' to *PI*; a reference to *PN* that assigns the value True implicitly assigns the value 'T' to *PI*.
- iv) If the caller language of *EP* is FORTRAN, then a reference to *PN* that assigns the value False to *PN* implicitly assigns the value .FALSE. to *PI*; a reference to *PN* that assigns the value True implicitly assigns the value .TRUE. to *PI*.
- v) If the caller language of *EP* is *PLI*, then a reference to *PN* that assigns the value False to *PN* implicitly assigns the value '0'B to *PI*; a reference to *PN* that assigns the value True implicitly assigns the value '1'B to *PI*.

NOTE 373 — Pascal has a Boolean-type, whose values are True and False.

## 13.4 Calls to an &lt;externally-invoked procedure&gt;

- g) If *P* is a binary large object locator parameter, a character large object locator parameter, an array locator parameter, a multiset locator parameter, or a user-defined type locator parameter, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the binary large object string locator value, the character large object locator value, the array locator value, the multiset locator value, or the user-defined type locator value, respectively, that uniquely identifies *SV* to *PI*.
- h) If *DT* identifies a CHARACTER LARGE OBJECT or BINARY LARGE OBJECT type, then

Case:

- i) If the caller language of *EP* is C, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value OCTET\_LENGTH(*SV*) to *PN.PN\_length* and the value *SV* to *PN.PN\_data*.
- ii) If the caller language of *EP* is COBOL, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value OCTET\_LENGTH(*SV*) to *PN.PN-OCTET\_LENGTH* and the value *SV* to *PN.PN-DATA*.
- iii) If the caller language of *EP* is FORTRAN, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value OCTET\_LENGTH(*SV*) to *PN\_LENGTH* and the value *SV* to *PN\_DATA*.
- iv) If the caller language of *EP* is PLI, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value OCTET\_LENGTH(*SV*) to *PN.PN\_length* and the value *SV* to *PN.PN\_data*.
- i) If *DT* identifies a BINARY(*L*) data type and the caller language of *EP* is C, then let *CL* be the maximum possible length in octets of *PN*. A reference to *PN* that assigns some value *SV* to *PN* implicitly assigns a value that is an SQL BINARY(*CL*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with X'00's as necessary to reach the length *CL*.
- j) If *DT* identifies a BINARY VARYING(*L*) data type and the caller language of *EP* is C, then a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value OCTET\_LENGTH(*SV*) to *PN.PN\_length* and the value *SV* to *PN.PN\_data*.
- k) If *DT* identifies a BINARY(*L*) data type and the caller language of *EP* is either COBOL or FORTRAN, then let *CL* be the maximum possible length in octets of *PN*. A reference to *PN* that assigns some value *SV* to *PN* implicitly assigns a value that is an SQL BINARY(*CL*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with X'00's as necessary to reach the length *CL*.
- l) If *DT* identifies a BINARY(*L*) or BINARY VARYING(*L*) data types and the caller language of *EP* is PLI, then let *CL* be the maximum possible length in octets of *PN*. A reference to *PN* that assigns some value *SV* to *PN* implicitly assigns a value that is

Case:

- i) If *DT* identifies BINARY(*L*), then an SQL BINARY(*CL*) data type.
- ii) Otherwise, an SQL BINARY VARYING(*CL*) data type in which octets of the value are the corresponding octets of *SV*, padded on the right with X'00's as necessary to reach the length *CL*.
- m) Otherwise, a reference to *PN* that assigns some value *SV* to *PN* implicitly assigns the value *SV* to *PI*. If the caller language of *EP* is ADA and no value has been assigned to *PI*, then an implementation-dependent value is assigned to *PI*.

## **Conformance Rules**

*None.*



## 13.5 <SQL procedure statement>

*This Subclause is modified by Subclause 11.2, “<SQL procedure statement>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 14.4, “<SQL procedure statement>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 13.2, “<SQL procedure statement>”, in ISO/IEC 9075-14.*

### Function

Define all of the SQL-statements that are <SQL procedure statement>s.

### Format

```
<SQL procedure statement> ::=
 <SQL executable statement>

<SQL executable statement> ::=
 <SQL schema statement>
 | <SQL data statement>
 | <SQL control statement>
 | <SQL transaction statement>
 | <SQL connection statement>
 | <SQL session statement>
 | <SQL diagnostics statement>
 | <SQL dynamic statement>

<SQL schema statement> ::=
 <SQL schema definition statement>
 | <SQL schema manipulation statement>

04 09 <SQL schema definition statement> ::=
 <schema definition>
 | <table definition>
 | <view definition>
 | <SQL-invoked routine>
 | <grant statement>
 | <role definition>
 | <domain definition>
 | <character set definition>
 | <collation definition>
 | <transliteration definition>
 | <assertion definition>
 | <trigger definition>
 | <user-defined type definition>
 | <user-defined cast definition>
 | <user-defined ordering definition>
 | <transform definition>
 | <sequence generator definition>

04 09 <SQL schema manipulation statement> ::=
 <drop schema statement>
 | <alter table statement>
 | <drop table statement>
 | <drop view statement>
 | <alter routine statement>
```

## 13.5 &lt;SQL procedure statement&gt;

```

| <drop routine statement>
| <drop user-defined cast statement>
| <revoke statement>
| <drop role statement>
| <alter domain statement>
| <drop domain statement>
| <drop character set statement>
| <drop collation statement>
| <drop transliteration statement>
| <drop assertion statement>
| <drop trigger statement>
| <alter type statement>
| <drop data type statement>
| <drop user-defined ordering statement>
| <alter transform statement>
| <drop transform statement>
| <alter sequence generator statement>
| <drop sequence generator statement>

<SQL data statement> ::=
| <open statement>
| <fetch statement>
| <close statement>
| <select statement: single row>
| <free locator statement>
| <hold locator statement>
| <SQL data change statement>

<SQL data change statement> ::=
| <delete statement: positioned>
| <delete statement: searched>
| <insert statement>
| <update statement: positioned>
| <update statement: searched>
| <truncate table statement>
| <merge statement>

04 <SQL control statement> ::=
| <call statement>
| <return statement>

<SQL transaction statement> ::=
| <start transaction statement>
| <set transaction statement>
| <set constraints mode statement>
| <savepoint statement>
| <release savepoint statement>
| <commit statement>
| <rollback statement>

<SQL connection statement> ::=
| <connect statement>
| <set connection statement>
| <disconnect statement>

09 14 <SQL session statement> ::=
| <set session user identifier statement>
| <set role statement>

```

```
| <set local time zone statement>
| <set session characteristics statement>
| <set catalog statement>
| <set schema statement>
| <set names statement>
| <set path statement>
| <set transform group statement>
| <set session collation statement>
```

```
04 <SQL diagnostics statement> ::=
 <get diagnostics statement>
```

```
<SQL dynamic statement> ::=
 <SQL descriptor statement>
| <prepare statement>
| <deallocate prepared statement>
| <describe statement>
| <execute statement>
| <execute immediate statement>
| <SQL dynamic data statement>
```

```
<SQL dynamic data statement> ::=
 <allocate cursor statement>
| <dynamic open statement>
| <dynamic fetch statement>
| <dynamic close statement>
| <dynamic delete statement: positioned>
| <dynamic update statement: positioned>
```

```
<SQL descriptor statement> ::=
 <allocate descriptor statement>
| <deallocate descriptor statement>
| <set descriptor statement>
| <get descriptor statement>
```

## Syntax Rules

- 1) Let *S* be the <SQL procedure statement>.
- 2) 04 An <SQL connection statement> shall not be generally contained in an <SQL control statement>.
- 3) The SQL-invoked routine specified by <SQL-invoked routine> shall be a schema-level routine.  
NOTE 374 — “schema-level routine” is defined in Subclause 11.51, “<SQL-invoked routine>”.
- 4) *S* is *possibly non-deterministic* if and only if *S* is not an <SQL schema statement> and at least one of the following is satisfied:
  - a) *S* is a <select statement: single row> that is possibly non-deterministic.
  - b) *S* contains a <routine invocation> whose subject routine is an SQL-invoked routine that is possibly non-deterministic.
  - c) *S* generally contains a <query specification> or a <query expression> that is possibly non-deterministic.
  - d) 04 *S* generally contains a <value expression> that is possibly non-deterministic.

## Access Rules

*None.*

## General Rules

- 1) Let *S* be the executing statement specified in an application of this Subclause.

NOTE 375 — *S* is necessarily the innermost executing statement of the SQL-session as defined in Subclause 4.37, “SQL-sessions”.

NOTE 376 — *S* is not necessarily an <SQL procedure statement>.

- 2) A *statement execution context NEWSEC* is established for the execution of *S*. Let *OLDSEC* be the most recent statement execution context. *NEWSEC* becomes the most recent statement execution context. *NEWSEC* is an atomic execution context, and therefore the most recent atomic execution context, if and only if *S* is an atomic SQL-statement.
- 3) The current trigger execution context *CTEC*, if any, is preserved and a new trigger execution context *NTEC* is created with an empty set of state changes *SSC*.
- 4) If the non-dynamic or dynamic execution of an <SQL data statement>, <SQL dynamic data statement>, <dynamic select statement>, or <dynamic single row select statement> occurs within the same SQL-transaction as the non-dynamic or dynamic execution of an SQL-schema statement and this is not allowed by the SQL-implementation, then an exception condition is raised: *invalid transaction state — schema and data statement mixing not supported*.
- 5) Case:
  - a) If *S* is immediately contained in an <externally-invoked procedure> *EP*, then let *n* be the number of <host parameter declaration>s specified in *EP*; let *PD<sub>i</sub>*, 1 (one) ≤ *i* ≤ *n*, be the *i*-th such <host parameter declaration>; and let *PN<sub>i</sub>* and *DT<sub>i</sub>* be the <host parameter name> and <data type>, respectively, specified in *PD<sub>i</sub>*. When *EP* is called by an SQL-agent, let *PI<sub>i</sub>* be the *i*-th argument in the procedure call.

Case:

- i) If *S* is an <SQL connection statement>, then:
  - 1) The SQL-client module that contains *S* is associated with the SQL-agent.
  - 2) The first diagnostics area is emptied.
  - 3) For each *i*, 1 (one) ≤ *i* ≤ *n*, the General Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, are evaluated for input parameters with *EP*, *PD<sub>i</sub>*, *PN<sub>i</sub>*, *DT<sub>i</sub>*, and *PI<sub>i</sub>* as *PROC*, *DECL*, *NAME*, *TYPE*, and *ARG*, respectively.
  - 4) The General Rules of *S* are evaluated.
  - 5) For each *i*, 1 (one) ≤ *i* ≤ *n*, the General Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, are evaluated for output parameters with *EP*, *PD<sub>i</sub>*, *PN<sub>i</sub>*, *DT<sub>i</sub>*, and *PI<sub>i</sub>* as *PROC*, *DECL*, *NAME*, *TYPE*, and *ARG*, respectively.
  - 6) If *S* successfully initiated or resumed an SQL-session, then subsequent calls to an <externally-invoked procedure> and subsequent invocations of <direct SQL statement>s by the SQL-

agent are associated with that SQL-session until the SQL-agent terminates the SQL-session or makes it dormant.

ii) If  $S$  is an <SQL diagnostics statement>, then:

- 1) The SQL-client module that contains  $S$  is associated with the SQL-agent.
- 2) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the General Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, are evaluated for input parameters with  $EP$ ,  $PD_i$ ,  $PN_i$ ,  $DT_i$ , and  $PI_i$  as  $PROC$ ,  $DECL$ ,  $NAME$ ,  $TYPE$ , and  $ARG$ , respectively.
- 3) The General Rules of  $S$  are evaluated.
- 4) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the General Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, are evaluated for output parameters with  $EP$ ,  $PD_i$ ,  $PN_i$ ,  $DT_i$ , and  $PI_i$  as  $PROC$ ,  $DECL$ ,  $NAME$ ,  $TYPE$ , and  $ARG$ , respectively.

iii) Otherwise:

- 1) If no SQL-session is current for the SQL-agent, then

Case:

- A) If the SQL-agent has not executed an <SQL connection statement> and there is no default SQL-session associated with the SQL-agent, then the following <connect statement> is effectively executed:

CONNECT TO DEFAULT

- B) If the SQL-agent has not executed an <SQL connection statement> and there is a default SQL-session associated with the SQL-agent, then the following <set connection statement> is effectively executed:

SET CONNECTION DEFAULT

- C) Otherwise, an exception condition is raised: *connection exception* — *connection does not exist*, no further subrules of this General Rule are evaluated, and evaluation continues with the next General Rule.
- 2) Subsequent calls to an <externally-invoked procedure> and subsequent invocations of <direct SQL statement>s by the SQL-agent are associated with the SQL-session until the SQL-agent terminates the SQL-session or makes it dormant.
- 3) If no SQL-transaction is active for the SQL-agent and  $S$  is an SQL-statement that implicitly initiates SQL-transactions, then an SQL-transaction is initiated. The SQL-client module that contains  $S$  is associated with the SQL-transaction.
- 4) The SQL-client module that contains  $S$  is associated with the SQL-agent.
- 5) If  $S$  contains an <SQL schema statement> and the access mode of the current SQL-transaction is read-only, then an exception condition is raised: *invalid transaction state* no further subrules of this General Rule are evaluated, and evaluation continues with the next General Rule.
- 6) 04 The first diagnostics area is emptied.

## 13.5 &lt;SQL procedure statement&gt;

- 7) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the General Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, are evaluated for input parameters with  $EP$ ,  $PD_i$ ,  $PN_i$ ,  $DT_i$ , and  $PI_i$  as  $PROC$ ,  $DECL$ ,  $NAME$ ,  $TYPE$ , and  $ARG$ , respectively.
  - 8) If  $S$  does not conform to the Syntax Rules and Access Rules of an <SQL procedure statement>, then an exception condition is raised: *syntax error or access rule violation* no further subrules of this General Rule are evaluated, and evaluation continues with the next General Rule.
  - 9) The General Rules of  $S$  are evaluated.
  - 10) For each  $i$ ,  $1 \text{ (one)} \leq i \leq n$ , the General Rules of Subclause 13.4, “Calls to an <externally-invoked procedure>”, are evaluated for output parameters with  $EP$ ,  $PD_i$ ,  $PN_i$ ,  $DT_i$ , and  $PI_i$  as  $PROC$ ,  $DECL$ ,  $NAME$ ,  $TYPE$ , and  $ARG$ , respectively.
  - 11) If  $S$  is a <select statement: single row> or a <fetch statement> and a completion condition is raised: *no data*, or an exception condition is raised, then the value of each  $PI_i$  for which  $PN_i$  is referenced in a <target specification> in  $S$  is implementation-dependent.
- b) Otherwise:
- i) If no SQL-transaction is active for the SQL-agent and  $S$  is an SQL-statement that implicitly initiates SQL-transactions, then an SQL-transaction is initiated.
  - ii) If  $S$  is not an <SQL diagnostics statement>, then the first diagnostics area is emptied.
  - iii) If  $S$  is an <SQL schema statement> and the access mode of the current SQL-transaction is read-only, then an exception condition is raised: *invalid transaction state*.
- 6) For every referential constraint  $RC$ , the General Rules of Subclause 15.17, “Execution of referential actions”, are applied with  $RC$  as *CONSTRAINT*.
  - 7) For every state change  $DSC$  in  $SSC$  whose trigger event is DELETE, let  $DSOT$  be the set of transitions in  $DSC$  and let  $DBT$  be the subject table of  $DSC$ . Each row that is marked for deletion from  $DBT$  is deleted from  $DBT$ .
  - 8) Case:
    - a) If, for some constraint  $C$ , the constraint mode of  $C$  in the current SQL-session is immediate and an applicable <search condition> included in the descriptor of  $C$  evaluates to *False*, then an exception condition is raised: *integrity constraint violation*.
 

NOTE 377 — This is needed to cater for both referential constraints specifying NO ACTION and referential constraints that are not satisfied after application of referential actions.
    - b) Otherwise, the Syntax Rules and General Rules of Subclause 15.18, “Execution of AFTER triggers”, are applied with  $SSC$  as the *SET OF STATE CHANGES*.
  - 9)  $NTEC$ , together with all of its contents, is destroyed and  $CTEC$ , if present, is restored to become the current trigger execution context.
  - 10) Case:
    - a) If  $S$  is immediately contained in an <externally-invoked procedure>, then
 

Case:

- i) If *S* executed successfully, then either a completion condition is raised: *successful completion*, or a completion condition is raised: *warning*, or a completion condition is raised: *no data*, as determined by the General Rules in this and other Subclauses of ISO/IEC 9075.
- ii) If *S* did not execute successfully, then:
  - 1) The status parameter is set to the value specified for the condition in [Clause 24](#), “Status codes”.
  - 2) If *S* is an atomic SQL-statement, then all changes made to SQL-data or schemas by the execution of *S* are canceled.

NOTE 378 — Atomic and non-atomic SQL-statements are defined in [Subclause 4.33.5](#), “SQL-statement atomicity and statement execution contexts”.

- b) Otherwise, the General Rules for *S* are evaluated.

Case:

- i) If *S* executed successfully, then either a completion condition is raised: *successful completion*, or a completion condition is raised: *warning*, or a completion condition is raised: *no data*, as determined by the General Rules in this and other Subclauses of ISO/IEC 9075.
- ii) Otherwise:
  - 1) If *S* is an atomic SQL-statement, then all changes made to SQL-data or schemas by the execution of *S* are canceled.

NOTE 379 — Atomic and non-atomic SQL-statements are defined in [Subclause 4.33.5](#), “SQL-statement atomicity and statement execution contexts”.

- 2) The exception condition with which the execution of *S* completed is raised.

- 11) If *S* is not an <SQL diagnostics statement>, then diagnostics information resulting from the execution of *S* is placed into the first diagnostics area, causing the first condition area in the first diagnostics area to become occupied. Whether any other condition areas become occupied is implementation-defined.
- 12) 04 If *NEWSEC* is atomic, then all savepoints established during its existence are destroyed.
- 13) *NEWSEC* ceases to exist and *OLDSEC* becomes the most recent statement execution context.
- 14) 04 *S* ceases to be an executing statement.

NOTE 380 — The innermost executing statement, if any, is now the one that was the innermost executing statement that caused *S* to be executed.

## Conformance Rules

*None.*

## 13.6 Data type correspondences

This Subclause is modified by Subclause 14.5, “Data type correspondences”, in ISO/IEC 9075-9.

This Subclause is modified by Subclause 13.3, “Data type correspondences”, in ISO/IEC 9075-14.

### Function

Specify the data type correspondences for SQL data types and host language types.

NOTE 381 — These tables are referenced in Subclause 11.51, “<SQL-invoked routine>”, for the definitions of external routines and in Subclause 10.4, “<routine invocation>”, for the invocation of external routines.

In the following tables, let  $P$  be <precision>,  $S$  be <scale>,  $L$  be <length> or <large object length>,  $U$  be the <char length units>,  $CS$  be the <character set specification>,  $T$  be <time fractional seconds precision>,  $Q$  be <interval qualifier>, and  $N$  be the implementation-defined size of a structured type reference.

### Tables

**Table 16 — Data type correspondences for Ada**

| SQL Data Type                                           | Ada Data Type                               |
|---------------------------------------------------------|---------------------------------------------|
| SQLSTATE                                                | SQL_STANDARD.SQLSTATE_TYPE                  |
| CHARACTER ( $L$ $U$ )<br>CHARACTER SET $CS$             | SQL_STANDARD.CHAR, with P'LENGTH of $L*k^1$ |
| CHARACTER VARYING ( $L$ )<br>CHARACTER SET $CS$         | <i>None</i>                                 |
| CHARACTER LARGE OBJECT<br>( $L$ )<br>CHARACTER SET $CS$ | <i>None</i>                                 |
| BINARY                                                  | <i>None</i>                                 |
| BINARY VARYING                                          | <i>None</i>                                 |
| BINARY LARGE OBJECT( $L$ )                              | <i>None</i>                                 |
| NUMERIC( $P,S$ )                                        | <i>None</i>                                 |
| DECIMAL( $P,S$ )                                        | <i>None</i>                                 |
| SMALLINT                                                | SQL_STANDARD.SMALLINT                       |
| INTEGER                                                 | SQL_STANDARD.INT                            |
| BIGINT                                                  | SQL_STANDARD.BIGINT                         |



| SQL Data Type                                                                                                                                                    | Ada Data Type                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| FLOAT( <i>P</i> )                                                                                                                                                | <i>None</i>                                 |
| REAL                                                                                                                                                             | SQL_STANDARD.REAL                           |
| DOUBLE PRECISION                                                                                                                                                 | SQL_STANDARD.DOUBLE_PRECISION               |
| BOOLEAN                                                                                                                                                          | SQL_STANDARD.BOOLEAN                        |
| DATE                                                                                                                                                             | <i>None</i>                                 |
| TIME( <i>T</i> )                                                                                                                                                 | <i>None</i>                                 |
| TIMESTAMP( <i>T</i> )                                                                                                                                            | <i>None</i>                                 |
| INTERVAL( <i>Q</i> )                                                                                                                                             | <i>None</i>                                 |
| user-defined type                                                                                                                                                | <i>None</i>                                 |
| REF                                                                                                                                                              | SQL_STANDARD.CHAR with P'LENGTH of <i>N</i> |
| ROW                                                                                                                                                              | <i>None</i>                                 |
| ARRAY                                                                                                                                                            | <i>None</i>                                 |
| MULTISET                                                                                                                                                         | <i>None</i>                                 |
| <sup>1</sup> If <i>U</i> is OCTETS, then <i>k</i> is 1 (one); Otherwise, <i>k</i> is the maximum number of octets per character in the character set <i>CS</i> . |                                             |

**Table 17 — Data type correspondences for C**

| SQL Data Type                                                     | C Data Type                                                                                                        |
|-------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| SQLSTATE                                                          | char, with length 6                                                                                                |
| CHARACTER ( <i>L U</i> )<br>CHARACTER SET <i>CS</i>               | <i>unit</i> , with length $(L+1)*k^1$                                                                              |
| CHARACTER VARYING ( <i>L U</i> )<br>CHARACTER SET <i>CS</i>       | <i>unit</i> , with length $(L+1)*k^1$                                                                              |
| CHARACTER LARGE OBJECT<br>( <i>L</i> )<br>CHARACTER SET <i>CS</i> | <pre> struct {     long hvn_reserved     unsigned long hvn_length     unit hvn_data[L]; } hvn<sup>1 2</sup> </pre> |
| BINARY ( <i>L</i> )                                               | char, with length <i>L</i>                                                                                         |

| SQL Data Type                    | C Data Type                                                                                                    |
|----------------------------------|----------------------------------------------------------------------------------------------------------------|
| BINARY VARYING ( <i>L</i> )      | <pre>struct {     long hvn_reserved     unsigned long hvn_length     char hvn_data[L]; } hvn<sup>2</sup></pre> |
| BINARY LARGE OBJECT ( <i>L</i> ) | <pre>struct {     long hvn_reserved     unsigned long hvn_length     char hvn_data[L]; } hvn<sup>2</sup></pre> |
| NUMERIC ( <i>P,S</i> )           | <i>None</i>                                                                                                    |
| DECIMAL ( <i>P,S</i> )           | <i>None</i>                                                                                                    |
| SMALLINT                         | pointer to short                                                                                               |
| INTEGER                          | pointer to long                                                                                                |
| BIGINT                           | pointer to long long                                                                                           |
| FLOAT ( <i>P</i> )               | <i>None</i>                                                                                                    |
| REAL                             | pointer to float                                                                                               |
| DOUBLE PRECISION                 | pointer to double                                                                                              |
| BOOLEAN                          | pointer to long                                                                                                |
| DATE                             | <i>None</i>                                                                                                    |
| TIME ( <i>T</i> )                | <i>None</i>                                                                                                    |
| TIMESTAMP ( <i>T</i> )           | <i>None</i>                                                                                                    |
| INTERVAL ( <i>Q</i> )            | <i>None</i>                                                                                                    |
| user-defined type                | <i>None</i>                                                                                                    |
| REF                              | char, with length <i>N</i>                                                                                     |
| ROW                              | <i>None</i>                                                                                                    |
| ARRAY                            | <i>None</i>                                                                                                    |
| MULTISET                         | <i>None</i>                                                                                                    |

<sup>1</sup> If *U* is OCTETS, then *unit* is char and *k* is 1 (one). Otherwise (that is, if *U* is CHARACTERS) then *unit* is an appropriate implementation-defined C data type (typically char, unsigned char, unsigned short, or unsigned int) and *k* is

| SQL Data Type                                                                                                                                                                                                                                                                                                                                                                                                                                                     | C Data Type |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| <p>an appropriate sizing factor given the choice of <i>unit</i>. The choice of <i>unit</i> and <i>k</i> is implementation-defined, based on the character set <i>CS</i>. For example, for UTF32, <i>unit</i> might be <code>char</code> and <i>k</i> might be 4; or <i>unit</i> might be <code>unsigned int</code> and <i>k</i> might be 1 (one).</p> <p><sup>2</sup> <i>hvn</i> is the name of the host variable defined to correspond to the SQL data type.</p> |             |

**Table 18 — Data type correspondences for COBOL**

| SQL Data Type                                                       | COBOL Data Type                                                                                                                                                                     |
|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SQLSTATE                                                            | PICTURE X(5)                                                                                                                                                                        |
| CHARACTER ( <i>L U</i> )<br>CHARACTER SET <i>CS</i>                 | PICTURE X( <i>L</i> ) <sup>3</sup>                                                                                                                                                  |
| CHARACTER VARYING ( <i>L U</i> )<br>CHARACTER SET <i>CS</i>         | <i>None</i>                                                                                                                                                                         |
| CHARACTER LARGE OBJECT<br>( <i>L U</i> )<br>CHARACTER SET <i>CS</i> | 01 <i>hvn</i> .<br>49 <i>hvn</i> -RESERVED PIC S9(9) USAGE IS BINARY.<br>49 <i>hvn</i> -LENGTH PIC S9(9) USAGE IS BINARY.<br>49 <i>hvn</i> -DATA PIC X( <i>L</i> ) <sup>2 3</sup> . |
| BINARY ( <i>L</i> )                                                 | PICTURE X( <i>L</i> )                                                                                                                                                               |
| BINARY VARYING ( <i>L</i> )                                         | <i>None</i>                                                                                                                                                                         |
| BINARY LARGE OBJECT ( <i>L</i> )                                    | 01 <i>hvn</i> .<br>49 <i>hvn</i> -RESERVED PIC S9(9) USAGE IS BINARY.<br>49 <i>hvn</i> -LENGTH PIC S9(9) USAGE IS BINARY.<br>49 <i>hvn</i> -DATA PIC X( <i>L</i> ) <sup>2</sup> .   |
| NUMERIC ( <i>P,S</i> )                                              | USAGE DISPLAY SIGN LEADING SEPARATE, with PICTURE as specified <sup>1</sup>                                                                                                         |
| DECIMAL( <i>P,S</i> )                                               | <i>None</i>                                                                                                                                                                         |
| SMALLINT                                                            | PICTURE S9( <i>SPI</i> ) USAGE BINARY, where <i>SPI</i> is implementation-defined                                                                                                   |
| INTEGER                                                             | PICTURE S9( <i>PI</i> ) USAGE BINARY, where <i>PI</i> is implementation-defined                                                                                                     |
| BIGINT                                                              | PICTURE S9( <i>BPI</i> ) USAGE BINARY, where <i>BPI</i> is implementation-defined                                                                                                   |
| FLOAT ( <i>P</i> )                                                  | <i>None</i>                                                                                                                                                                         |
| REAL                                                                | <i>None</i>                                                                                                                                                                         |

| SQL Data Type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | COBOL Data Type                   |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------|
| DOUBLE PRECISION                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <i>None</i>                       |
| BOOLEAN                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | PICTURE X                         |
| DATE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <i>None</i>                       |
| TIME (T)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <i>None</i>                       |
| TIMESTAMP (T)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <i>None</i>                       |
| INTERVAL (Q)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <i>None</i>                       |
| user-defined type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <i>None</i>                       |
| REF                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | alphanumeric with length <i>N</i> |
| ROW                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <i>None</i>                       |
| ARRAY                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <i>None</i>                       |
| MULTISET                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <i>None</i>                       |
| <p><sup>1</sup> Case:</p> <p>a) If <math>S=P</math>, then a PICTURE with an 'S' followed by a 'V' followed by <math>P</math> '9's.</p> <p>b) If <math>P &gt; S &gt; 0</math>, then a PICTURE with an 'S' followed by <math>P-S</math> '9's followed by a 'V' followed by <math>S</math> '9's.</p> <p>c) If <math>S=0</math>, then a PICTURE with an 'S' followed by <math>P</math> '9's optionally followed by a 'V'.</p> <p><sup>2</sup> <i>hvn</i> is the name of the host variable defined to correspond to the SQL data type</p> <p><sup>3</sup> If <math>U</math> is CHARACTERS, then for character set UTF16, as well as other implementation-defined character sets in which a code unit occupies two octets, "PICTURE X(<math>L</math>)" should be replaced with "PICTURE N(<math>L</math>)". For character set UTF32, as well as other implementation-defined character sets in which a code unit occupies four octets, "PICTURE X(<math>L</math>)" should be replaced with "PICTURE (????)". Otherwise, "PICTURE X(<math>L</math>)" should be used.</p> |                                   |

**09 14 Table 19 — Data type correspondences for Fortran**

| SQL Data Type                                       | Fortran Data Type                        |
|-----------------------------------------------------|------------------------------------------|
| SQLSTATE                                            | CHARACTER, with length 5                 |
| CHARACTER ( $L$ $U$ )<br>CHARACTER SET $CS$         | CHARACTER <sup>2</sup> , with length $L$ |
| CHARACTER VARYING ( $L$ $U$ )<br>CHARACTER SET $CS$ | <i>None</i>                              |

| SQL Data Type                                                       | Fortran Data Type                                                                                                                                                                                                                                     |
|---------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHARACTER LARGE OBJECT<br>( <i>L U</i> )<br>CHARACTER SET <i>CS</i> | CHARACTER <i>hvn</i> ( <i>L</i> +8)<br>INTEGER*4 <i>hvn_RESERVED</i><br>INTEGER*4 <i>hvn_LENGTH</i><br>CHARACTER <i>hvn_DATA</i><br>EQUIVALENCE( <i>hvn</i> (5), <i>hvn_LENGTH</i> )<br>EQUIVALENCE( <i>hvn</i> (9), <i>hvn_DATA</i> ) <sup>1 2</sup> |
| BINARY ( <i>L</i> )                                                 | CHARACTER, with length <i>L</i>                                                                                                                                                                                                                       |
| BINARY VARYING ( <i>L</i> )                                         | <i>None</i>                                                                                                                                                                                                                                           |
| BINARY LARGE OBJECT ( <i>L</i> )                                    | CHARACTER <i>hvn</i> ( <i>L</i> +8)<br>INTEGER*4 <i>hvn_RESERVED</i><br>INTEGER*4 <i>hvn_LENGTH</i><br>CHARACTER <i>hvn_DATA</i><br>EQUIVALENCE( <i>hvn</i> (5), <i>hvn_LENGTH</i> )<br>EQUIVALENCE( <i>hvn</i> (9), <i>hvn_DATA</i> ) <sup>1</sup>   |
| NUMERIC ( <i>P,S</i> )                                              | <i>None</i>                                                                                                                                                                                                                                           |
| DECIMAL ( <i>P,S</i> )                                              | <i>None</i>                                                                                                                                                                                                                                           |
| SMALLINT                                                            | <i>None</i>                                                                                                                                                                                                                                           |
| INTEGER                                                             | INTEGER                                                                                                                                                                                                                                               |
| BIGINT                                                              | <i>None</i>                                                                                                                                                                                                                                           |
| FLOAT ( <i>P</i> )                                                  | <i>None</i>                                                                                                                                                                                                                                           |
| REAL                                                                | REAL                                                                                                                                                                                                                                                  |
| DOUBLE PRECISION                                                    | DOUBLE PRECISION                                                                                                                                                                                                                                      |
| BOOLEAN                                                             | LOGICAL                                                                                                                                                                                                                                               |
| DATE                                                                | <i>None</i>                                                                                                                                                                                                                                           |
| TIME ( <i>T</i> )                                                   | <i>None</i>                                                                                                                                                                                                                                           |
| TIMESTAMP ( <i>T</i> )                                              | <i>None</i>                                                                                                                                                                                                                                           |
| INTERVAL ( <i>Q</i> )                                               | <i>None</i>                                                                                                                                                                                                                                           |
| user-defined type                                                   | <i>None</i>                                                                                                                                                                                                                                           |
| REF                                                                 | CHARACTER with length <i>N</i>                                                                                                                                                                                                                        |
| ROW                                                                 | <i>None</i>                                                                                                                                                                                                                                           |

| SQL Data Type                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Fortran Data Type |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| ARRAY                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <i>None</i>       |
| MULTISET                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <i>None</i>       |
| <sup>1</sup> <i>hvn</i> is the name of the host variable defined to correspond to the SQL data type<br><sup>2</sup> If <i>U</i> is CHARACTERS, then for character set UTF16, as well as other implementation-defined character sets in which a code unit occupies more than one octet, “CHARACTER KIND= <i>n</i> ” should be used; in this case, the value of <i>n</i> that corresponds to a given character set is implementation-defined. Otherwise, “CHARACTER” (without “KIND= <i>n</i> ”) should be used. |                   |

**Table 20 — Data type correspondences for M**

| SQL Data Type                       | M Data Type                               |
|-------------------------------------|-------------------------------------------|
| SQLSTATE                            | character, with maximum length at least 5 |
| CHARACTER ( <i>L</i> )              | <i>None</i>                               |
| CHARACTER VARYING ( <i>L</i> )      | character with maximum length <i>L</i>    |
| CHARACTER LARGE OBJECT ( <i>L</i> ) | <i>None</i>                               |
| BINARY ( <i>L</i> )                 | <i>None</i>                               |
| BINARY VARYING ( <i>L</i> )         | <i>None</i>                               |
| BINARY LARGE OBJECT ( <i>L</i> )    | <i>None</i>                               |
| NUMERIC ( <i>P,S</i> )              | character                                 |
| DECIMAL ( <i>P,S</i> )              | character                                 |
| SMALLINT                            | <i>None</i>                               |
| INTEGER                             | character                                 |
| BIGINT                              | <i>None</i>                               |
| FLOAT ( <i>P</i> )                  | <i>None</i>                               |
| REAL                                | character                                 |
| DOUBLE PRECISION                    | <i>None</i>                               |
| BOOLEAN                             | <i>None</i>                               |
| DATE                                | <i>None</i>                               |
| TIME ( <i>T</i> )                   | <i>None</i>                               |

| SQL Data Type          | M Data Type |
|------------------------|-------------|
| TIMESTAMP ( <i>T</i> ) | <i>None</i> |
| INTERVAL ( <i>Q</i> )  | <i>None</i> |
| user-defined type      | <i>None</i> |
| REF                    | character   |
| ROW                    | <i>None</i> |
| ARRAY                  | <i>None</i> |
| MULTISET               | <i>None</i> |

**09 14 Table 21 — Data type correspondences for Pascal**

| SQL Data Type                        | Pascal Data Type                     |
|--------------------------------------|--------------------------------------|
| SQLSTATE                             | PACKED ARRAY [1..5] OF CHAR          |
| CHARACTER (1)                        | CHAR                                 |
| CHARACTER ( <i>L</i> ), <i>L</i> > 1 | PACKED ARRAY [1.. <i>L</i> ] OF CHAR |
| CHARACTER VARYING ( <i>L</i> )       | <i>None</i>                          |
| CHARACTER LARGE OBJECT ( <i>L</i> )  | <i>None</i>                          |
| BINARY ( <i>L</i> )                  | <i>None</i>                          |
| BINARY VARYING ( <i>L</i> )          | <i>None</i>                          |
| BINARY LARGE OBJECT ( <i>L</i> )     | <i>None</i>                          |
| NUMERIC ( <i>P,S</i> )               | <i>None</i>                          |
| DECIMAL ( <i>P,S</i> )               | <i>None</i>                          |
| SMALLINT                             | <i>None</i>                          |
| INTEGER                              | INTEGER                              |
| BIGINT                               | <i>None</i>                          |
| FLOAT ( <i>P</i> )                   | <i>None</i>                          |
| REAL                                 | REAL                                 |

| SQL Data Type     | Pascal Data Type           |
|-------------------|----------------------------|
| DOUBLE PRECISION  | <i>None</i>                |
| BOOLEAN           | BOOLEAN                    |
| DATE              | <i>None</i>                |
| TIME (T)          | <i>None</i>                |
| TIMESTAMP (T)     | <i>None</i>                |
| INTERVAL (Q)      | <i>None</i>                |
| user-defined type | <i>None</i>                |
| REF               | PACKED ARRAY[1..N] OF CHAR |
| ROW               | <i>None</i>                |
| ARRAY             | <i>None</i>                |
| MULTISET          | <i>None</i>                |

**Table 22 — Data type correspondences for PL/I**

| SQL Data Type              | PL/I Data Type                                                                                                            |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------|
| SQLSTATE                   | CHARACTER(5)                                                                                                              |
| CHARACTER (L)              | CHARACTER(L)                                                                                                              |
| CHARACTER VARYING (L)      | CHARACTER (L) VARYING                                                                                                     |
| CHARACTER LARGE OBJECT (L) | DCL 01 hvn<br>49 hvn_reserved FIXED BINARY (31)<br>49 hvn_length FIXED BINARY (31)<br>49 hvn_data CHAR (n) <sup>1</sup> ; |
| BINARY (L)                 | CHARACTER (L)                                                                                                             |
| BINARY VARYING (L)         | CHARACTER (L) VARYING                                                                                                     |
| BINARY LARGE OBJECT (L)    | DCL 01 hvn<br>49 hvn_reserved FIXED BINARY (31)<br>49 hvn_length FIXED BINARY (31)<br>49 hvn_data CHAR (n) <sup>1</sup> ; |
| NUMERIC(P,S)               | <i>None</i>                                                                                                               |



| SQL Data Type                                                                                       | PL/I Data Type                                                         |
|-----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| DECIMAL ( <i>P,S</i> )                                                                              | FIXED DECIMAL ( <i>P,S</i> )                                           |
| SMALLINT                                                                                            | FIXED BINARY( <i>SPI</i> ), where <i>SPI</i> is implementation-defined |
| INTEGER                                                                                             | FIXED BINARY( <i>PI</i> ), where <i>PI</i> is implementation-defined   |
| BIGINT                                                                                              | FIXED BINARY( <i>BPI</i> ), where <i>BPI</i> is implementation-defined |
| FLOAT ( <i>P</i> )                                                                                  | FLOAT BINARY ( <i>P</i> )                                              |
| REAL                                                                                                | <i>None</i>                                                            |
| DOUBLE PRECISION                                                                                    | <i>None</i>                                                            |
| BOOLEAN                                                                                             | BIT(1)                                                                 |
| DATE                                                                                                | <i>None</i>                                                            |
| TIME ( <i>T</i> )                                                                                   | <i>None</i>                                                            |
| TIMESTAMP ( <i>T</i> )                                                                              | <i>None</i>                                                            |
| INTERVAL ( <i>Q</i> )                                                                               | <i>None</i>                                                            |
| user-defined type                                                                                   | <i>None</i>                                                            |
| REF                                                                                                 | CHARACTER( <i>N</i> ) VARYING                                          |
| ROW                                                                                                 | <i>None</i>                                                            |
| ARRAY                                                                                               | <i>None</i>                                                            |
| MULTISET                                                                                            | <i>None</i>                                                            |
| <sup>1</sup> <i>hvn</i> is the name of the host variable defined to correspond to the SQL data type |                                                                        |

## Conformance Rules

*None.*

*(Blank page)*

## 14 Data manipulation

### 14.1 <declare cursor>

*This Subclause is modified by Subclause 12.1, “<declare cursor>”, in ISO/IEC 9075-4.*

#### Function

Declare a standing cursor.

#### Format

```
<declare cursor> ::=
 DECLARE <cursor name> <cursor properties>
 FOR <cursor specification>
```

#### Syntax Rules

- 1) 04 If a <declare cursor> is contained in an <SQL-client module definition> *M*, then:
  - a) The <cursor name> shall not be equivalent to the <cursor name> of any other <declare cursor> or <dynamic declare cursor> in *M*.
  - b) The scope of the <cursor name> is *M* with the exception of any <SQL schema statement> contained in *M*.
  - c) Any <host parameter name> contained in the <cursor specification> shall be defined in a <host parameter declaration> in the <externally-invoked procedure> that contains an <open statement> that specifies the <cursor name> and is contained in the scope of that <cursor name>.

NOTE 382 — See the Syntax Rules of Subclause 13.1, “<SQL-client module definition>”.

#### Access Rules

*None.*

#### General Rules

- 1) A cursor declaration descriptor *CDD* is created. *CDD* includes indications that:
  - a) The kind of cursor is a standing cursor.
  - b) 04 The provenance of the cursor is an indication of the SQL-client module whose <SQL-client module definition> contains the <declare cursor>.

**14.1 <declare cursor>**

- c) The name of the cursor is the <cursor name>.
- d) The cursor's origin is the <cursor specification> contained in the <declare cursor>.
- e) The cursor's declared properties are as determined by the <cursor properties>.

**Conformance Rules**

- 1) Without Feature F831, “Full cursor update”, conforming SQL language shall not contain a <declare cursor> that contains both a <cursor specification> that contains an <updatability clause> that specifies FOR UPDATE and <cursor properties> that contain a <cursor scrollability>.

## 14.2 <cursor properties>

### Function

Specify the declared properties of a cursor.

### Format

```
<cursor properties> ::=
 [<cursor sensitivity>] [<cursor scrollability>] CURSOR
 [<cursor holdability>]
 [<cursor returnability>]

<cursor sensitivity> ::=
 SENSITIVE
 | INSENSITIVE
 | ASENSITIVE

<cursor scrollability> ::=
 SCROLL
 | NO SCROLL

<cursor holdability> ::=
 WITH HOLD
 | WITHOUT HOLD

<cursor returnability> ::=
 WITH RETURN
 | WITHOUT RETURN
```

### Syntax Rules

- 1) If <cursor sensitivity> is not specified, then ASENSITIVE is implicit.
- 2) If <cursor scrollability> is not specified, then NO SCROLL is implicit.
- 3) If <cursor holdability> is not specified, then WITHOUT HOLD is implicit.
- 4) If <cursor returnability> is not specified, then WITHOUT RETURN is implicit.

### Access Rules

*None.*

### General Rules

- 1) The declared properties of the cursor declaration descriptor associated with <cursor properties> are given by:
  - a) The declared sensitivity property is the explicit or implicit <cursor sensitivity>.
  - b) The declared scrollability property is the explicit or implicit <cursor scrollability>.

- c) The declared holdability property is the explicit or implicit <cursor holdability>.
- d) The declared returnability property is the explicit or implicit <cursor returnability>.

## **Conformance Rules**

- 1) Without Feature T231, “Sensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains SENSITIVE.
- 2) Without Feature F791, “Insensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains INSENSITIVE.
- 3) Without Feature F791, “Insensitive cursors”, or Feature T231, “Sensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains ASENSITIVE.
- 4) Without Feature F431, “Read-only scrollable cursors”, conforming SQL language shall not contain a <cursor scrollability>.
- 5) Without Feature T471, “Result sets return value”, conforming SQL language shall not contain a <cursor returnability>.
- 6) Without Feature T551, “Optional key words for default syntax”, conforming SQL language shall not contain a <cursor holdability> that immediately contains WITHOUT HOLD.

## 14.3 <cursor specification>

### Function

Define a result set.

### Format

```
<cursor specification> ::=
 <query expression> [<updatability clause>]

<updatability clause> ::=
 FOR { READ ONLY | UPDATE [OF <column name list>] }
```

### Syntax Rules

- 1) Let *CS* be the <cursor specification>.
- 2) Let *QE* be the <query expression> immediately contained in *CS*.
- 3) Case:
  - a) If *CS* is contained in a <declare cursor> *DC* that contains a <cursor properties>, then let *CP* be that <cursor properties>.
  - b) If *CS* is a <preparable statement> being prepared by a <prepare statement> or re-prepared by an <allocate cursor statement> or a <dynamic open statement>, then:
    - i) Let *STMT* be the <prepare statement> that is preparing *CS*, or the <allocate cursor statement> or <dynamic open statement> that is re-preparing *CS*. Let *SCMD* be the <SQL-client module definition> that contains *STMT*.
    - ii) Case:
      - 1) If *CS* is being prepared by a <prepare statement>, then let *PS* be that <prepare statement>.
      - 2) Otherwise, let *PS* be the <prepare statement> that previously prepared *CS*.
    - iii) Let *SSV* be the <SQL statement variable> contained in *PS*.
    - iv) Case:
      - 1) If *CS* is being prepared by a <prepare statement>, then:  
Case:
        - A) If *SSV* is a <statement name> and there is exactly one <dynamic declare cursor> *DDC* contained in *SCMD* whose <statement name> is equivalent to *SSV*, then let *CPI* be the <cursor properties> contained in *DDC*.
        - B) If *SSV* is an <extended statement name> that specifies or implies a <scope option> that is *LOCAL*, and there is exactly one <allocate cursor statement> *ACS* contained in *SCMD* whose <cursor intent> is <statement cursor> and whose <extended statement name> specifies or implies *LOCAL*, then let *CPI* be the <cursor properties> contained in *ACS*.

- C) Otherwise, let *CP1* be a zero-length string.
- 2) If *CS* is being re-prepared by an <allocate cursor statement> *ACS*, then let *CP1* be the <cursor properties> contained in *ACS*.
- 3) If *CS* is being re-prepared by a <dynamic open statement>, then let *DDC* be the <dynamic declare cursor> whose <cursor name> is equivalent to the <cursor name> contained in *STMT*. Let *CP1* be the <cursor properties> contained in *DDC*.
- v) If *PS* contains an <attributes variable>, then let *CP2* be the value of that <attributes variable>; otherwise, let *CP2* be the zero-length string.
- vi) Case:
  - 1) If *CP2* contains <cursor sensitivity>, then let *SENS* be that <cursor sensitivity>.
  - 2) If *CP1* contains <cursor sensitivity>, then let *SENS* be that <cursor sensitivity>.
  - 3) Otherwise, let *SENS* be the zero-length string.
- vii) Case:
  - 1) If *CP2* contains <cursor scrollability>, then let *SCRO* be that <cursor scrollability>.
  - 2) If *CP1* contains <cursor scrollability>, then let *SCRO* be that <cursor scrollability>.
  - 3) Otherwise, let *SCRO* be the zero-length string.
- viii) Case:
  - 1) If *CP2* contains <cursor holdability>, then let *HOLD* be that <cursor holdability>.
  - 2) If *CP1* contains <cursor holdability>, then let *HOLD* be that <cursor holdability>.
  - 3) Otherwise, let *HOLD* be the zero-length string.
- ix) Case:
  - 1) If *CP2* contains <cursor returnability>, then let *RET* be that <cursor returnability>.
  - 2) If *CP1* contains <cursor returnability>, then let *RET* be that <cursor returnability>.
  - 3) Otherwise, let *RET* be the zero-length string.
- x) Let *CP* be the <cursor properties>:

*SENS SCRO CURSOR HOLD RET*

- 4) If <updatability clause> is not specified and either *CS* is contained in a <declare cursor> or is being re-prepared by an <allocate cursor statement> or a <dynamic open statement>, then

Case:

- a) If *CP* contains *INSENSITIVE* or *SCROLL*, or *QE* immediately contains an <order by clause>, or *QE* is not a simply updatable <query specification>, then an <updatability clause> of *READ ONLY* is implicit.
- b) Otherwise, an <updatability clause> of *FOR UPDATE* without a <column name list> is implicit.



NOTE 383 — If *CS* is being prepared by a <prepare statement>, then defaulting the <updatability clause> is postponed until *CS* is re-prepared.

- 5) If an <updatability clause> of FOR UPDATE with or without a <column name list> is specified, then *CP* shall not contain INSENSITIVE, *QE* shall be updatable, and *QE* shall have only one leaf underlying table *LUT* such that *QE* is one-to-one with respect to *LUT*.
- 6) Case:
  - a) If an <updatability clause> specifying FOR UPDATE is specified or implicit, then *CS* is *updatable*.
  - b) If an <updatability clause> specifying FOR READ ONLY is specified or implicit, then *CS* is *not updatable*.
  - c) Otherwise, the determination of updatability of *CS* is postponed until *CS* is re-prepared.
- 7) If *CS* is updatable, then let *LUTN* be a <table name> that references *LUT*. *LUTN* is an exposed <table or query name> whose scope is <updatability clause>.
- 8) If an <updatability clause> of FOR UPDATE without a <column name list> is specified or implicit, then a <column name list> that consists of the <column name> of every column of *LUT* is implicit.
- 9) If an <updatability clause> of FOR UPDATE with a <column name list> is specified, then each <column name> in the <column name list> shall be the <column name> of a column of *LUT*.

## Access Rules

*None.*

## General Rules

*None.*

## Conformance Rules

- 1) Without Feature F831, “Full cursor update”, conforming SQL language shall not contain a <cursor specification> that contains both an <updatability clause> that specifies FOR UPDATE and an <order by clause>.
- 2) Without Feature T111, “Updatable joins, unions, and columns”, in conforming SQL language, if FOR UPDATE is specified, then *QE* shall be simply updatable.

## 14.4 <open statement>

*This Subclause is modified by Subclause 12.2, “<open statement>”, in ISO/IEC 9075-4.*

### Function

Open a standing cursor.

### Format

```
<open statement> ::=
 OPEN <cursor name>
```

### Syntax Rules

- 1) 04 Let *CN* be the <cursor name> in the <open statement>. *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*.
- 2) *CN* shall identify a standing cursor.
- 3) Let *CDD* be the cursor declaration descriptor of the standing cursor identified specified by *CN*.

### Access Rules

- 1) The Access Rules for the <query expression> simply contained in the <declare cursor> identified by the <cursor name> are applied.

### General Rules

- 1) Let *CR* be the cursor instance descriptor in the current SQL-session whose cursor declaration descriptor is *CDD*.
- 2) The General Rules of Subclause 15.1, “Effect of opening a cursor”, are applied with *CR* as *CURS*OR.

### Conformance Rules

*None.*

## 14.5 <fetch statement>

*This Subclause is modified by Subclause 12.3, “<fetch statement>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 11.15, “<fetch statement>”, in ISO/IEC 9075-10.  
This Subclause is modified by Subclause 14.1, “<fetch statement>”, in ISO/IEC 9075-14.*

### Function

Position a standing cursor on a specified row of the standing cursor's result set and retrieve values from that row.

### Format

```

10 <fetch statement> ::=
 FETCH [[<fetch orientation>] FROM] <cursor name> INTO <fetch target list>

<fetch orientation> ::=
 NEXT
 | PRIOR
 | FIRST
 | LAST
 | { ABSOLUTE | RELATIVE } <simple value specification>

14 <fetch target list> ::=
 <target specification> [{ <comma> <target specification> }...]

```

### Syntax Rules

- 1) <fetch target list> shall not contain a <target specification> that specifies a <column reference>.
- 2) If the <fetch orientation> is omitted, then NEXT is implicit.
- 3) 04 10 Let *CN* be the <cursor name> in the <fetch statement>. *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*.
- 4) *CN* shall identify a standing cursor.
- 5) Let *CDD* be the cursor declaration descriptor of the standing cursor identified by *CN*.
- 6) Let *T* be the result set defined by the <cursor specification> of *CDD*.
- 7) If the implicit or explicit <fetch orientation> is not NEXT, then the declared scrollability property of *CDD* shall be SCROLL.
- 8) If a <fetch orientation> that contains a <simple value specification> is specified, then the declared type of that <simple value specification> shall be exact numeric with a scale of 0 (zero).
- 9) Case:
  - a) If the <fetch target list> contains a single <target specification> *TS* and the degree of *T* is greater than 1 (one), then the declared type of *TS* shall be a row type.
 Case:

- i) 04 If  $TS$  is an <SQL parameter reference>, then the Syntax Rules of Subclause 9.2, “Store assignment”, apply to  $TS$  as *TARGET* and the row type of  $T$  as *VALUE*.
  - ii) Otherwise, the Syntax Rules of Subclause 9.1, “Retrieval assignment”, apply to  $TS$  as *TARGET* and the row type of  $T$  as *VALUE*.
- b) Otherwise:
- i) The number of <target specification>s  $NTS$  in the <fetch target list> shall be the same as the degree of  $T$ . The  $i$ -th <target specification>,  $1 \text{ (one)} \leq i, \leq NTS$ , in the <fetch target list> corresponds with the  $i$ -th column of  $T$ .
  - ii) 04 For  $i$  varying from 1 (one) to  $NTS$ , let  $TSI_i$  be the  $i$ -th <target specification> in the <fetch target list> that is either an <SQL parameter reference> or a <target array element specification>, and let  $CS_i$  be the  $i$ -th column of  $T$  that corresponds with the <target specification> in the <fetch target list>.
- Case:
- 1) If  $TSI_i$  contains a <simple value specification>, then the Syntax Rules of Subclause 9.2, “Store assignment”, apply to an arbitrary site whose declared type is the declared type of  $TSI_i$  and  $CS_i$  as *TARGET* and *VALUE*, respectively.
  - 2) Otherwise, the Syntax Rules of Subclause 9.2, “Store assignment”, apply to  $TSI_i$  and the corresponding column of  $T$  as *TARGET* and *VALUE*, respectively.
- iii) 10 For each <target specification>  $TS2_i$ ,  $1 \text{ (one)} \leq i, \leq NTS$ , that is a <host parameter specification>, the Syntax Rules of Subclause 9.1, “Retrieval assignment”, apply to  $TS2_i$  as *TARGET* and the corresponding column of  $T$  as *VALUE*.
  - iv) For each <target specification>  $TS2_i$ ,  $1 \text{ (one)} \leq i, \leq NTS$ , that is an <embedded variable specification>, the Syntax Rules of Subclause 9.1, “Retrieval assignment”, apply to  $TS2_i$  as *TARGET* and the corresponding column of  $T$  as *VALUE*.

## Access Rules

*None.*

## General Rules

- 1) Let  $CR$  be the cursor instance descriptor of the current SQL-session whose cursor declaration descriptor is  $CDD$ .
- 2) If  $CR$  is not in the open state, then an exception condition is raised: *invalid cursor state*.
- 3) The General Rules of Subclause 15.3, “Determination of the current row of a cursor”, are applied with  $CR$  as *CURSOR* and <fetch orientation> as *FETCH ORIENTATION*.
- 4) If a completion condition *no data* has been raised, then no further General Rules of this Subclause are applied.
- 5) Case:

- a) If the <fetch target list> contains a single <target specification> *TS* and the degree of *T* is greater than 1 (one), then the current row is assigned to *TS* and

Case:

- i) 04 14 If *TS* is an <SQL parameter reference>, then the General Rules of Subclause 9.2, “Store assignment”, are applied to *TS* and the current row as *TARGET* and *VALUE*, respectively.
  - ii) Otherwise, the General Rules of Subclause 9.1, “Retrieval assignment”, are applied to *TS* and the current row as *TARGET* and *VALUE*, respectively.
- b) 10 Otherwise, if the <fetch target list> contains more than one <target specification>, then values from the current row are assigned to their corresponding targets identified by the <fetch target list>. The assignments are made in an implementation-dependent order. Let *TV* be a target and let *SV* denote its corresponding value in the current row of *CR*.

Case:

- i) 04 If *TV* is either an <SQL parameter reference> or a <target array element specification>, then for each <target specification> in the <fetch target list>, let *TV<sub>i</sub>* be the *i*-th <target specification> in the <fetch target list> and let *SV<sub>i</sub>* denote the *i*-th corresponding value in the current row of *CR*.

Case:

- 1) If <target array element specification> is specified, then

Case:

- A) If the value of *TV<sub>i</sub>*, denoted by *C*, is the null value, then an exception condition is raised: *data exception — null value in array target*.

- B) Otherwise:

- I) Let *N* be the maximum cardinality of *C*.
- II) Let *M* be the cardinality of the value of *C*.
- III) Let *I* be the value of the <simple value specification> immediately contained in *TV<sub>i</sub>*.
- IV) Let *EDT* be the element type of *C*.
- V) Case:

- 1) If *I* is greater than zero and less than or equal to *M*, then the value of *C* is replaced by an array *A* with element type *EDT* and cardinality *M* derived as follows:
  - a) For *j* varying from 1 (one) to *I*–1 and from *I*+1 to *M*, the *j*-th element in *A* is the value of the *j*-th element in *C*.
  - b) 14 The *I*-th element of *A* is set to the value of *SV*, denoted by *SV<sub>i</sub>*, by applying the General Rules of Subclause 9.2, “Store assignment”, to the *I*-th element of *A* and *SV<sub>i</sub>* as *TARGET* and *VALUE*, respectively.

- 2) If  $I$  is greater than  $M$  and less than or equal to  $N$ , then the value of  $C$  is replaced by an array  $A$  with element type  $EDT$  and cardinality  $I$  derived as follows:
    - a) For  $j$  varying from 1 (one) to  $M$ , the  $j$ -th element in  $A$  is the value of the  $j$ -th element in  $C$ .
    - b) For  $j$  varying from  $M+1$  to  $I$ , the  $j$ -th element in  $A$  is the null value.
    - c) 14 The  $I$ -th element of  $A$  is set to the value of  $SV$ , denoted by  $SV_i$ , by applying the General Rules of Subclause 9.2, “Store assignment”, to the  $I$ -th element of  $A$  and  $SV_i$  as *TARGET* and *VALUE*, respectively.
  - 3) Otherwise, an exception condition is raised: *data exception — array element error*.
- 2) 14 Otherwise, the General Rules of Subclause 9.2, “Store assignment”, are applied to  $TV_i$  and  $SV_i$  as *TARGET* and *VALUE*, respectively.
  - ii) If  $TV$  is a <host parameter name>, then the General Rules of Subclause 9.1, “Retrieval assignment”, are applied to  $TV$  and  $SV$  as *TARGET* and *VALUE*, respectively.
  - iii) If  $TV$  is an <embedded variable specification>, then the General Rules of Subclause 9.1, “Retrieval assignment”, are applied to  $TV$  and  $SV$  as *TARGET* and *VALUE*, respectively.

NOTE 384 — SQL parameters cannot have as their data types any row type.

- 6) If an exception condition occurs during the assignment of a value to a target, then the values of all targets are implementation-dependent.

NOTE 385 — It is implementation-dependent whether  $CR$  remains positioned on the current row when an exception condition is raised during the derivation of any <derived column>.

## Conformance Rules

- 1) Without Feature F431, “Read-only scrollable cursors”, a <fetch statement> shall not contain a <fetch orientation>.

## 14.6 <close statement>

*This Subclause is modified by Subclause 12.4, “<close statement>”, in ISO/IEC 9075-4.*

### Function

Close a standing cursor.

### Format

```
<close statement> ::=
 CLOSE <cursor name>
```

### Syntax Rules

- 1) 04 Let *CN* be the <cursor name> in the <close statement>. *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*.
- 2) *CN* shall identify a standing cursor.
- 3) Let *CDD* be the cursor declaration descriptor of the standing cursor identified by *CN*.

### Access Rules

*None.*

### General Rules

- 1) Let *CR* be the cursor instance descriptor of the current SQL-session whose cursor declaration descriptor is *CDD*.
- 2) The General Rules of Subclause 15.4, “Effect of closing a cursor”, are applied with *CR* as *CURSOR* and DESTROY as *DISPOSITION*.

### Conformance Rules

*None.*

## 14.7 <select statement: single row>

*This Subclause is modified by Subclause 12.5, “<select statement: single row>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 11.14, “<select statement: single row>”, in ISO/IEC 9075-10.*

*This Subclause is modified by Subclause 14.2, “<select statement: single row>”, in ISO/IEC 9075-14.*

### Function

Retrieve values from a specified row of a table.

### Format

```
<select statement: single row> ::=
 SELECT [<set quantifier>] <select list>
 INTO <select target list>
 <table expression>
```

```
14 <select target list> ::=
 <target specification> [{ <comma> <target specification> }...]
```

### Syntax Rules

- 1) <select target list> shall not contain a <target specification> that specifies a <column reference>.
- 2) Let *T* be the table defined by the <table expression>.
- 3) Case:
  - a) If the <select target list> contains a single <target specification> *TS* and the degree of table *T* is greater than 1 (one), then the declared type of *TS* shall be a row type.
 

Case:

    - i) If *TS* is an <SQL parameter reference>, then the Syntax Rules of Subclause 9.2, “Store assignment”, apply to *TS* and the row type of table *T* as *TARGET* and *VALUE*, respectively.
    - ii) Otherwise, the Syntax Rules of Subclause 9.1, “Retrieval assignment”, apply to *TS* and an arbitrary value of the row type of table *T* as *TARGET* and *VALUE*, respectively.
  - b) Otherwise:
    - i) The number of elements *NOE* in the <select list> shall be the same as the number of elements in the <select target list>. The *i*-th <target specification>,  $1 \text{ (one)} \leq i \leq NOE$ , in the <select target list> corresponds with the *i*-th element of the <select list>.
    - ii) 04 For *i* varying from 1 (one) to *NOE*, let *TS<sub>i</sub>* be the *i*-th <target specification> in the <select target list> that is either an <SQL parameter reference> or a <target array element specification>, and let *SL<sub>i</sub>* be the *i*-th element of the <select list> that corresponds to the <target specification> in the <select target list>.

Case:



- 1) If <target array element specification> is specified, then the Syntax Rules of Subclause 9.2, “Store assignment”, apply to an arbitrary site whose declared type is the declared type of  $TS_i$  and  $SL_i$  as *TARGET* and *VALUE*, respectively.
  - 2) Otherwise, the Syntax Rules of Subclause 9.2, “Store assignment”, apply to  $TS_i$  and the corresponding element of the <select list>, as *TARGET* and *VALUE*, respectively.
  - iii) For each <target specification>  $TS$  that is a <host parameter specification>, the Syntax Rules of Subclause 9.1, “Retrieval assignment”, apply to  $TS$  and the corresponding element of the <select list>, as *TARGET* and *VALUE*, respectively.
  - iv) 10 For each <target specification>  $TS$  that is an <embedded variable specification>, the Syntax Rules of Subclause 9.1, “Retrieval assignment”, apply to  $TS$  and the corresponding element of the <select list>, as *TARGET* and *VALUE*, respectively.
- 4) Let  $S$  be a <query specification> whose <select list> and <table expression> are those specified in the <select statement: single row> and that specifies the <set quantifier> if it is specified in the <select statement: single row>.  $S$  shall be a valid <query specification>.
  - 5) A column in the result of the <select statement: single row> is known not null if the corresponding column in the result of  $S$  is known not null.
  - 6) The <select statement: single row> is *possibly non-deterministic* if  $S$  is possibly non-deterministic.

## Access Rules

*None.*

## General Rules

- 1) Let  $Q$  be the result of <query specification>  $S$ .
- 2) Case:
  - a) If the cardinality of  $Q$  is greater than 1 (one), then an exception condition is raised: *cardinality violation*. It is implementation-dependent whether or not SQL-data values are assigned to the targets identified by the <select target list>.
  - b) If  $Q$  is empty, then no SQL-data values are assigned to any targets identified by the <select target list>, and a completion condition is raised: *no data*.
  - c) Otherwise, values in the row of  $Q$  are assigned to their corresponding targets.
- 3) If a completion condition *no data* has been raised, then no further General Rules of this Subclause are applied.
- 4) Case:
  - a) If the <select target list> contains a single <target specification>  $TS$  and the degree of table  $T$  is greater than 1 (one), then the current row is assigned to  $TS$  and  
Case:

- i) 14 If  $TS$  is an <SQL parameter reference>, then the General Rules of Subclause 9.2, “Store assignment”, apply to  $TS$  and the current row as  $TARGET$  and  $VALUE$ , respectively.
  - ii) Otherwise, the General Rules of Subclause 9.1, “Retrieval assignment”, are applied to  $TS$  and the current row as  $TARGET$  and  $VALUE$ , respectively.
- b) Otherwise:
- i) Let  $NOE$  be the number of elements in the <select list>.
  - ii) 04 For  $i$  varying from 1 (one) to  $NOE$ , let  $TS_i$  be the  $i$ -th <target specification> in the <select target list> that is either an <SQL parameter reference> or a <target array element specification>, and let  $SL_i$  denote the corresponding ( $i$ -th) value in the row of  $Q$ . The assignment of values to targets in the <select target list> is in an implementation-dependent order.

Case:

- 1) If <target array element specification> is specified, then

Case:

- A) If the value of  $TS_i$ , denoted by  $C$ , is the null value, then an exception condition is raised:  
*data exception — null value in array target.*

B) Otherwise:

- I) Let  $N$  be the maximum cardinality of  $C$ .
- II) Let  $M$  be the cardinality of the value of  $C$ .
- III) Let  $I$  be the value of the <simple value specification> immediately contained in  $TS_i$ .
- IV) Let  $EDT$  be the element type of  $C$ .
- V) Case:
  - 1) If  $I$  is greater than zero and less than or equal to  $M$ , then the value of  $C$  is replaced by an array  $A$  with element type  $EDT$  and cardinality  $M$  derived as follows:
    - a) For  $j$  varying from 1 (one) to  $I-1$  and from  $I+1$  to  $M$ , the  $j$ -th element in  $A$  is the value of the  $j$ -th element in  $C$ .
    - b) 14 The  $I$ -th element of  $A$  is set to the value of  $SL_i$ , denoted by  $SL_i$ , by applying the General Rules of Subclause 9.2, “Store assignment”, to the  $I$ -th element of  $A$  and  $SL_i$  as  $TARGET$  and  $VALUE$ , respectively.
  - 2) If  $I$  is greater than  $M$  and less than or equal to  $N$ , then the value of  $C$  is replaced by an array  $A$  with element type  $EDT$  and cardinality  $I$  derived as follows:
    - a) For  $j$  varying from 1 (one) to  $M$ , the  $j$ -th element in  $A$  is the value of the  $j$ -th element in  $C$ .
    - b) For  $j$  varying from  $M+1$  to  $I-1$ , the  $j$ -th element in  $A$  is the null value.

- c) 14 The  $I$ -th element of  $A$  is set to the value of  $SL$ , denoted by  $SL_i$ , by applying the General Rules of Subclause 9.2, “Store assignment”, to the  $I$ -th element of  $A$  and  $SL_i$  as *TARGET* and *VALUE*, respectively.
  - 3) Otherwise, an exception condition is raised: *data exception — array element error*.
- 2) 14 Otherwise, the corresponding value  $SL_i$  in the row of  $Q$  is assigned to  $TS_i$  according to the General Rules of Subclause 9.2, “Store assignment”, as *VALUE* and *TARGET*, respectively.
- iii) For each <target specification>  $TS$  that is a <host parameter specification>, the corresponding value in the row of  $Q$  is assigned to  $TS$  according to the General Rules of Subclause 9.1, “Retrieval assignment”, as *VALUE* and *TARGET*, respectively. The assignment of values to targets in the <select target list> is in an implementation-dependent order.
- iv) 10 For each <target specification>  $TS$  that is an <embedded variable specification>, the corresponding value in the row of  $Q$  is assigned to  $TS$  according to the General Rules of Subclause 9.1, “Retrieval assignment”, as *VALUE* and *TARGET*, respectively. The assignment of values to targets in the <select target list> is in an implementation-dependent order.
- 5) If an exception condition is raised during the assignment of a value to a target, then the values of all targets are implementation-dependent.

## Conformance Rules

*None.*

## 14.8 <delete statement: positioned>

This Subclause is modified by *Subclause 12.6*, “<delete statement: positioned>”, in ISO/IEC 9075-4.

This Subclause is modified by *Subclause 11.12*, “<delete statement: positioned>”, in ISO/IEC 9075-10.

### Function

Delete a row of a table.

### Format

```

10 <delete statement: positioned> ::=
 DELETE FROM <target table> [[AS] <correlation name>]
 WHERE CURRENT OF <cursor name>

<target table> ::=
 <table name>
 | ONLY <left paren> <table name> <right paren>

```

### Syntax Rules

- 1) 04 Let *DSP* be the <delete statement: positioned> and let *CN* be the <cursor name> immediately contained in *DSP*. *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*.
- 2) *CN* shall identify a standing cursor.
- 3) Let *CDD* be the cursor declaration descriptor of the standing cursor identified by *CN*.
- 4) The cursor specification of *CDD* shall be updatable.
- 5) Let *TN* be the <table name> contained in <target table>.
- 6) If <target table> *TT* immediately contains ONLY and the table identified by *TN* is not a typed table, then *TT* is equivalent to *TN*.
- 7) Let *T* be the simply underlying table of *CDD*. *T* is the *subject table* of the <delete statement: positioned>. Let *LUT* be the leaf underlying table of *T* such that *T* is one-to-one with respect to *LUT*.
- 8) The subject table of a <delete statement: positioned> shall not identify an old transition table or a new transition table.
- 9) *TN* shall identify *LUT*.
- 10) <target table> shall specify ONLY if and only if the <table reference> contained in *T* that references *LUT* specifies ONLY.
- 11) The schema identified by the explicit or implicit qualifier of *TN* shall include the descriptor of *LUT*.
- 12) Case:
  - a) If <correlation name> is specified, then let *CN* be that <correlation name>.
  - b) Otherwise, let *CN* be the <table name> contained in <target table>. *CN* is an exposed <table or query name>.

NOTE 386 — *CN* has no scope.

## Access Rules

- 1) Case:
  - a) If <delete statement: positioned> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges for the owner of that schema shall include DELETE for *TN*.
  - b) Otherwise, the current privileges shall include DELETE for *TN*.

## General Rules

- 1) Let *CR* be the cursor instance descriptor of the current SQL-session whose cursor declaration descriptor is *CDD*.
- 2) The General Rules of Subclause 15.5, “Effect of a positioned delete”, are applied with *CR* as *CURSOR*, *DSP* as *STATEMENT*, and *TT* as *TARGET*.

## Conformance Rules

- 1) Without Feature S111, “ONLY in query expressions”, conforming SQL language shall not contain a <target table> that contains ONLY.

## 14.9 <delete statement: searched>

This Subclause is modified by *Subclause 14.3*, “<delete statement: searched>”, in ISO/IEC 9075-14.

### Function

Delete rows of a table.

### Format

```
14 <delete statement: searched> ::=
 DELETE FROM <target table> [[AS] <correlation name>]
 [WHERE <search condition>]
```

### Syntax Rules

- 1) Let *TN* be the <table name> contained in the <target table>. Let *T* be the table identified by *TN*.
- 2) *T* shall be an updatable table or a trigger deletable table.
- 3) If the <delete statement: searched> is contained in a <triggered SQL statement>, then the <search condition> shall not contain a <value specification> that specifies a parameter reference.
- 4) *T* is the *subject table* of the <delete statement: searched>.
- 5) *TN* shall not identify an old transition table or a new transition table.
- 6) Case:
  - a) If <correlation name> is specified, then let *CN* be that <correlation name>.
  - b) Otherwise, let *CN* be the <table name> contained in <target table>. *CN* is an exposed <table or query name>.
- 7) The scope of *CN* is <search condition>.
- 8) If WHERE <search condition> is not specified, then WHERE TRUE is implicit.
- 9) The <search condition> shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.

### Access Rules

- 1) Case:
  - a) If <delete statement: searched> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema.
    - i) The applicable privileges for *A* shall include DELETE for *TN*.

- ii) If <target table> immediately contains ONLY, then the applicable privileges for A shall include SELECT WITH HIERARCHY OPTION on at least one supertable of T.
- b) Otherwise,
  - i) The current privileges shall include DELETE for TN.
  - ii) If <target table> immediately contains ONLY, then the current privileges shall include SELECT WITH HIERARCHY OPTION on at least one supertable of T.

## General Rules

- 1) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only, and T is not a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 2) If there is any sensitive cursor CR that is currently open in the SQL-transaction in which this SQL-statement is being executed, then

Case:

- a) If CR has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of this statement shall be made visible to CR or an exception condition is raised: *cursor sensitivity exception — request failed*.
- b) Otherwise, whether the change resulting from the successful execution of this SQL-statement is made visible to CR is implementation-defined.
- 3) If there is any open, insensitive cursor CR, then either the change resulting from the successful execution of this statement shall be invisible to CR, or an exception condition is raised: *cursor sensitivity exception — request failed*.
- 4) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- 5) The <search condition> is effectively evaluated for each row of T with the exposed <correlation name>s or <table or query name>s of the <table reference> bound to that row.

6) Case:

- a) If <target table> contains ONLY, then the rows for which the result of the <search condition> is True and for which there is no subrow in a proper subtable of T are identified for deletion from T.
- b) Otherwise, the rows for which the result of the <search condition> is True are identified for deletion from T.

NOTE 387 — Identifying a row for deletion is an implementation-dependent mechanism.

7) Case:

- a) If T is a base table, then:

i) Case:

- 1) If <target table> specifies ONLY, then T is *identified for deletion processing without subtables*.

- 2) Otherwise, *T* is *identified for deletion processing with subtables*.

NOTE 388 — Identifying a base table for deletion processing, with or without subtables, is an implementation-dependent mechanism.

- ii) The General Rules of Subclause 15.7, “Effect of deleting rows from base tables”, are applied.
- b) If *T* is a viewed table, then the General Rules of Subclause 15.9, “Effect of deleting some rows from a viewed table”, are applied with <target table> as *VIEW NAME*.
- 8) If any row that is marked for deletion by the <delete statement: searched> has been marked for deletion by any <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned> that identifies some open cursor *CR* or updated by any <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies some open cursor *CR*, then a completion condition is raised: *warning — cursor operation conflict*.
- 9) If no rows are marked for deletion, then a completion condition is raised: *no data*.

## Conformance Rules

- 1) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <delete statement: searched> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in the <search condition>.
- 2) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain a <delete statement: searched> that contains a <target table> that identifies a table that is not simply updatable.



## 14.10 <truncate table statement>

### Function

Delete all rows of a base table without causing any triggered action.

### Format

```
<truncate table statement> ::=
 TRUNCATE TABLE <target table>
```

### Syntax Rules

- 1) Let *TN* be the <table name> contained in the <target table>. Let *T* be the table identified by *TN*. The schema identified by the explicit or implicit <schema name> of *TN* shall include the descriptor of *T*.
- 2) *T* shall be a base table.
- 3) *T* shall not be identified by the name of the referenced table in any referential constraint descriptor.

### Access Rules

- 1) Let *A* be the <authorization identifier> that owns the schema identified by the <schema name> of the table identified by *TN*.
- 2) The enabled authorization identifiers shall include *A*.

### General Rules

- 1) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only, and *T* is not a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 2) If there is any sensitive cursor *CR* that is currently open in the SQL-transaction in which this SQL-statement is being executed, then  
Case:
  - a) If *CR* has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of this statement shall be made visible to *CR* or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - b) Otherwise, whether the change resulting from the successful execution of this SQL-statement is made visible to *CR* is implementation-defined.
- 3) If there is any open, insensitive cursor *CR*, then either the change resulting from the successful execution of this statement shall be invisible to *CR*, or an exception condition is raised: *cursor sensitivity exception — request failed*.
- 4) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.

- 5) Case:
  - a) If <target table> contains ONLY, then the rows for which there is no subrow in a proper subtable of *T* are deleted from *T*.
  - b) Otherwise, all rows are deleted from *T*.
- 6) If any row that is deleted from *T* by this <truncate table statement> has been marked for deletion by any <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned> that identifies some open cursor *CR* or updated by any <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies some open cursor *CR*, then a completion condition is raised: *warning — cursor operation conflict*.
- 7) If no rows are deleted from *T*, then a completion condition is raised: *no data*.

## Conformance Rules

- 1) Without Feature F200, “TRUNCATE TABLE statement”, conforming SQL language shall not contain a <truncate table statement>.

## 14.11 <insert statement>

*This Subclause is modified by Subclause 14.4, “<insert statement>”, in ISO/IEC 9075-14.*

### Function

Create new rows in a table.

### Format

```
14 <insert statement> ::=
 INSERT INTO <insertion target> <insert columns and source>

<insertion target> ::=
 <table name>

<insert columns and source> ::=
 <from subquery>
 | <from constructor>
 | <from default>

<from subquery> ::=
 [<left paren> <insert column list> <right paren>]
 [<override clause>]
 <query expression>

<from constructor> ::=
 [<left paren> <insert column list> <right paren>]
 [<override clause>]
 <contextually typed table value constructor>

<override clause> ::=
 OVERRIDING USER VALUE
 | OVERRIDING SYSTEM VALUE

<from default> ::=
 DEFAULT VALUES

<insert column list> ::=
 <column name list>
```

### Syntax Rules

- 1) Let *TN* be the <table name>; let *T* be the table identified by *TN*. If *T* is a view, then <target table> is effectively replaced by:  
  
ONLY ( *TN* )
- 2) *T* shall be insertable-into or trigger insertable-into.
- 3) For each leaf generally underlying table of *T* whose descriptor includes a user-defined type name *UDTN*, the data type descriptor of the user-defined type *UDT* identified by *UDTN* shall indicate that *UDT* is instantiable.

- 4) A column identified by the <insert column list> is an object column.
- 5) If *T* is not trigger insertable-into, then *T* shall be an updatable table; each object column of *T* shall be an updatable column.

NOTE 389 — The notion of updatable columns of base tables is defined in Subclause 4.14, “Tables”. The notion of updatable columns of viewed tables is defined in Subclause 11.23, “<view definition>”.

- 6) *T* is the *subject table* of the <insert statement>.
- 7) *TN* shall not identify an old transition table or a new transition table.
- 8) An <insert columns and source> that specifies DEFAULT VALUES is implicitly replaced by an <insert columns and source> that specifies a <contextually typed table value constructor> of the form

VALUES (DEFAULT, DEFAULT, . . . , DEFAULT)

where the number of “DEFAULT” entries is equal to the number of columns of *T*.

- 9) If *T* is not trigger insertable-into, then each <column name> in the <insert column list> shall identify an updatable column of *T*.
- 10) No <column name> of *T* shall be identified more than once. If the <insert column list> is omitted, then an <insert column list> that identifies all columns of *T* in the ascending sequence of their ordinal positions within *T* is implicit.
- 11) If <contextually typed table value constructor> *CTTVC* is specified, then every <contextually typed row value constructor element> simply contained in *CTTVC* whose positionally corresponding <column name> in <insert column list> references a column of which some underlying column is a generated column shall be a <default specification>.
- 12) Case:
  - a) If some underlying column of a column referenced by a <column name> contained in <insert column list> is a system-generated self-referencing column or a derived self-referencing column, then <override clause> shall be specified.
  - b) If for some *n*, some underlying column of the column referenced by the <column name> *CN* contained in the *n*-th ordinal position in <insert column list> is an identity column whose descriptor includes an indication that values are always generated, then

Case:

- i) If <from subquery> is specified, then <override clause> shall be specified.
- ii) If any <contextually typed row value expression> simply contained in the <contextually typed table value constructor> is a <row value special case>, then <override clause> shall be specified.
- iii) If the *n*-th <contextually typed row value constructor element> simply contained in any <contextually typed row value constructor> simply contained in the <contextually typed table value constructor> is not a <default specification>, then <override clause> shall be specified.

NOTE 390 — The preceding subrule does not cover all possibilities. The remaining possibilities are where <default clause> is specified for every identity column, in which case it is immaterial whether <override clause> is specified or not.

- c) If for some *n*, some underlying column of the column referenced by the <column name> *CN* contained in the *n*-th ordinal position in <insert column list> is an identity column whose descriptor includes an

indication that values are generated by default, then if <override clause> is specified, then <override clause> shall specify OVERRIDING USER VALUE.

- d) Otherwise, <override clause> shall not be specified.
- 13) If <contextually typed table value constructor> *CVC* is specified, then the data type of every <contextually typed value specification> *CVS* specified in every <contextually typed row value expression> *CRVS* contained in *CVC* is the data type *DT* indicated in the column descriptor for the positionally corresponding column in the explicit or implicit <insert column list>. If *CVS* is an <empty specification> that specifies ARRAY, then *DT* shall be an array type. If *CVS* is an <empty specification> that specifies MULTISSET, then *DT* shall be a multiset type.
- 14) Let *QT* be the table specified by the <query expression> or <contextually typed table value constructor>. The degree of *QT* shall be equal to the number of <column name>s in the <insert column list>. The column of table *T* identified by the *i*-th <column name> in the <insert column list> corresponds with the *i*-th column of *QT*.
- 15) The Syntax Rules of [Subclause 9.2, “Store assignment”](#), apply to corresponding columns of *T* and *QT* as *TARGET* and *VALUE*, respectively.
- 16) If the <insert statement> is contained in a <triggered SQL statement>, then the insert value shall not contain a <value specification> that specifies a parameter reference.
- 17) A <query expression> simply contained in a <from subquery> shall not be a <table value constructor>.

NOTE 391 — This rule removes a syntactic ambiguity; otherwise, “VALUES (1)” could be parsed either as

```
<insert columns and source> ::=
 <from subquery> ::=
 <query expression> ::=
 <table value constructor> ::=
 VALUES (1)
```

or

```
<insert columns and source> ::=
 <from constructor> ::=
 <contextually typed table value constructor> ::=
 VALUES (1)
```

## Access Rules

- 1) Case:
- a) If <insert statement> is contained in, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema. The applicable privileges for *A* for *TN* shall include INSERT for each object column.
- b) Otherwise, the current privileges for *TN* shall include INSERT for each object column.

## General Rules

- 1) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only, and *T* is not a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 2) If there is any sensitive cursor *CR* that is currently open in the SQL-transaction in which this SQL-statement is being executed, then

Case:

- a) If *CR* has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of this statement shall be made visible to *CR* or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - b) Otherwise, whether the change resulting from the successful execution of this SQL-statement is made visible to *CR* is implementation-defined.
- 3) If there is any open, insensitive cursor *CR*, then either the change resulting from the successful execution of this statement shall be invisible to *CR*, or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - 4) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
  - 5) *QT* is effectively evaluated before insertion of any rows into *T*.
  - 6) Let *Q* be the result of evaluating *QT*.
  - 7) For each row *R* of *Q*:
    - a) A candidate row of *T* is effectively created in which the value of each column is its default value, as specified in the General Rules of [Subclause 11.5](#), “<default clause>”. The candidate row consists of every column of *T*.
    - b) If *T* has a column *RC* of which some underlying column is a self-referencing column, then

Case:

- i) If *RC* is a system-generated self-referencing column, then the value of *RC* is effectively replaced by the REF value of the candidate row.
  - ii) If *RC* is a derived self-referencing column, then the value of *RC* is effectively replaced by a value derived from the columns in the candidate row that correspond to the list of attributes of the derived representation of the reference type of *RC* in an implementation-dependent manner.
- c) For each object column in the candidate row, let *C<sub>i</sub>* be the object column identified by the *i*-th <column name> in the <insert column list> and let *SV<sub>i</sub>* be the *i*-th value of *R*.
  - d) For every *C<sub>i</sub>* for which one of the following conditions is true:
    - i) *C<sub>i</sub>* is not marked as unassigned and no underlying column of *C<sub>i</sub>* is a self-referencing column.
    - ii) Some underlying column of *C<sub>i</sub>* is a user-generated self-referencing column.

- iii) Some underlying column of  $C_i$  is a self-referencing column and OVERRIDING SYSTEM VALUE is specified.
- iv) Some underlying column of  $C_i$  is an identity column and the  $i$ -th column of  $R$  is not derived from <default specification> and OVERRIDING SYSTEM VALUE is specified.
- v) Some underlying column of  $C_i$  is an identity column whose descriptor includes an indication that values are generated by default and neither OVERRIDING USER VALUE is specified nor is the  $i$ -th column derived from <default specification>.

the General Rules of Subclause 9.2, “Store assignment”, are applied with  $C_i$  and  $SV_i$  as *TARGET* and *SOURCE*, respectively.  $C_i$  is no longer marked as unassigned.

NOTE 392 — The data values allowable in the candidate row may be constrained by a WITH CHECK OPTION constraint. The effect of a WITH CHECK OPTION constraint is defined in the General Rules of Subclause 15.12, “Effect of inserting a table into a viewed table”.

- 8) Let  $S$  be the table consisting of the candidate rows.

Case:

- a) If  $T$  is a base table, then:

- i)  $T$  is identified for insertion of source table  $S$ .

NOTE 393 — Identifying a base table for insertion of a source table is an implementation-dependent operation.

- ii) The General Rules of Subclause 15.10, “Effect of inserting tables into base tables”, are applied.

- b) If  $T$  is a viewed table, then the General Rules of Subclause 15.12, “Effect of inserting a table into a viewed table”, are applied with  $S$  as *SOURCE* and  $T$  as *TARGET*.

- 9) If  $Q$  is empty, then a completion condition is raised: *no data*.

## Conformance Rules

- 1) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain an <insert statement> in which the <table name> of a leaf generally underlying table of  $T$  is generally contained in the <from subquery> except as the table name of a qualifying table of a column reference.
- 2) Without Feature F222, “INSERT statement: DEFAULT VALUES clause”, conforming SQL language shall not contain a <from default>.
- 3) Without Feature S024, “Enhanced structured types”, in conforming SQL language, for each column  $C$  identified in the explicit or implicit <insert column list>, if the declared type of  $C$  is a structured type  $TY$ , then the declared type of the corresponding column of the <query expression> or <contextually typed table value constructor> shall be  $TY$ .
- 4) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain an <override clause>.
- 5) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain an <insert statement> that contains an <insertion target> that identifies a table that is not simply updatable.

## 14.12 <merge statement>

*This Subclause is modified by Subclause 14.5, “<merge statement>”, in ISO/IEC 9075-14.*

### Function

Conditionally update rows of a table, or insert new rows into a table, or both.

### Format

```
14 <merge statement> ::=
 MERGE INTO <target table> [[AS] <merge correlation name>]
 USING <table reference>
 ON <search condition> <merge operation specification>

<merge correlation name> ::=
 <correlation name>

<merge operation specification> ::=
 <merge when clause>...

<merge when clause> ::=
 <merge when matched clause>
 | <merge when not matched clause>

<merge when matched clause> ::=
 WHEN MATCHED [AND <search condition>]
 THEN <merge update specification>

<merge when not matched clause> ::=
 WHEN NOT MATCHED [AND <search condition>]
 THEN <merge insert specification>

<merge update specification> ::=
 UPDATE SET <set clause list>

<merge insert specification> ::=
 INSERT [<left paren> <insert column list> <right paren>]
 [<override clause>]
 VALUES <merge insert value list>

<merge insert value list> ::=
 <left paren>
 <merge insert value element> [{ <comma> <merge insert value element> }...]
 <right paren>

<merge insert value element> ::=
 <value expression>
 | <contextually typed value specification>
```

### Syntax Rules

- 1) Let *TN* be the <table name> contained in <target table> and let *T* be the table identified by *TN*. *T* is the *subject table* of the <merge statement>.



- 2) Case:
  - a) If <merge when not matched clause> is specified, then *T* shall be insertable-into or trigger insertable-into.
  - b) Otherwise, *T* shall be updatable or trigger updatable.
- 3) *T* shall not be an old transition table or a new transition table.
- 4) For each leaf generally underlying table of *T* whose descriptor includes a user-defined type name *UDTN*, the data type descriptor of the user-defined type *UDT* identified by *UDTN* shall indicate that *UDT* is instantiable.
- 5) If *T* is a view, then <target table> is effectively replaced by:  
  
ONLY ( *TN* )
- 6) Case:
  - a) If <merge correlation name> is specified, then let *CN* be the <correlation name> contained in <merge correlation name>.
  - b) Otherwise, let *CN* be the <table name> contained in <target table>.
- 7) The scope of *CN* is the <search condition> immediately contained in the <merge statement>, the <search condition> immediately contained in a <merge when matched clause>, the <search condition> immediately contained in a <merge when not matched clause>, and the <set clause list>.
- 8) Let *TR* be the <table reference> immediately contained in <merge statement>. *TR* shall not directly contain a <joined table>.
- 9) The <correlation name> or exposed <table name> that is exposed by *TR* shall not be equivalent to *CN*.
- 10) If an <insert column list> is omitted, then an <insert column list> that identifies all columns of *T* in the ascending sequence of their ordinal position within *T* is implicit.
- 11) Case:
  - a) If some underlying column of a column referenced by a <column name> contained in <insert column list> is a system-generated self-referencing column or a derived self-referencing column, then <override clause> shall be specified.
  - b) If for some *n*, some underlying column of the column referenced by the <column name> *CN* contained in the *n*-th ordinal position in <insert column list> is an identity column whose descriptor includes an indication that values are always generated, and the *n*-th <contextually typed value specification> simply contained in any <merge insert value element> simply contained in the <merge insert value list> is not a <default specification>, then <override clause> shall be specified.
  - c) If for some *n*, some underlying column of the column referenced by the <column name> *CN* contained in the *n*-th ordinal position in <insert column list> is an identity column whose descriptor includes an indication that values are generated by default and <override clause> is specified, then <override clause> shall specify OVERRIDING USER VALUE.
  - d) Otherwise, <override clause> shall not be specified.
- 12) The <search condition> immediately contained in a <merge statement>, the <search condition> immediately contained in a <merge when matched clause>, and the <search condition> immediately contained in a

<merge when not matched clause> shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.

- 13) Each column identified by an <object column> in a <set clause list> is an *update object column*. Each column identified by a <column name> in an implicit or explicit <insert column list> is an *insert object column*. Each update object column and each insert object column is an *object column*.
- 14) If <merge when not matched clause> is specified and if *T* is not trigger insertable-into or if <merge when matched clause> is specified and if *T* is not trigger updatable, then every object column shall identify an updatable column of *T*.

NOTE 394 — The notion of updatable columns of base tables is defined in Subclause 4.14, “Tables”. The notion of updatable columns of viewed tables is defined in Subclause 11.23, “<view definition>”.

- 15) No <column name> of *T* shall be identified more than once in an <insert column list>.
- 16) For each <merge when not matched clause>:
- Let *NI* be the number of <merge insert value element>s contained in <merge insert value list>. Let *EXP*<sub>1</sub>, *EXP*<sub>2</sub>, ..., *EXP*<sub>*NI*</sub> be those <merge insert value element>s.
  - The number of <column name>s in the <insert column list> shall be equal to *NI*.
  - The declared type of every <contextually typed value specification> *CVS* in a <merge insert value list> is the data type *DT* indicated in the column descriptor for the positionally corresponding column in the explicit or implicit <insert column list>. If *CVS* is an <empty specification> that specifies ARRAY, then *DT* shall be an array type. If *CVS* is an <empty specification> that specifies MULTISSET, then *DT* shall be a multiset type.
  - Every <merge insert value element> whose positionally corresponding <column name> in <insert column list> references a column of which some underlying column is a generated column shall be a <default specification>.
  - For  $1 \text{ (one)} \leq i \leq NI$ , the Syntax Rules of Subclause 9.2, “Store assignment”, are applied with *EXP*<sub>*i*</sub> as *VALUE* and the column of table *T* identified by the *i*-th <column name> in the <insert column list> as *TARGET*.

## Access Rules

- 1) Case:
- If <merge statement> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema.
    - The applicable privileges for *A* shall include UPDATE for each update object column.
    - The applicable privileges for *A* shall include INSERT for each insert object column.
    - If <target table> immediately contains ONLY, then the applicable privileges for *A* shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *T*.
  - Otherwise,
    - The current privileges shall include UPDATE for each update object column.

- ii) The current privileges shall include INSERT for each insert object column.
- iii) If <target table> immediately contains ONLY, then the current privileges shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *T*.

## General Rules

- 1) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only, and *T* is not a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 2) If there is any sensitive cursor *CR* that is currently open in the SQL-transaction in which this SQL-statement is being executed, then

Case:

- a) If *CR* has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of this statement shall be made visible to *CR* or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - b) Otherwise, whether the change resulting from the successful execution of this SQL-statement is made visible to *CR* is implementation-defined.
- 3) If there is any open, insensitive cursor *CR*, then either the change resulting from the successful execution of this statement shall be invisible to *CR*, or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - 4) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
  - 5) Let *QT* be the table specified by the <table reference>. *QT* is effectively evaluated before update or insertion of any rows in *T*. Let *Q* be the result of evaluating *QT*.
  - 6) For each <merge when clause>, in the order specified in the <merge operation specification>,

Case:

- a) If <merge when matched clause> *MWMC* is specified, then:
  - i) For each row *RI* of *T*:
    - 1) The <search condition> immediately contained in <merge statement> and the <search condition> immediately contained in *MWMC*, if any, are effectively evaluated for *RI* with the exposed <table name> of the <target table> bound to *RI* and to each row of *Q* with the exposed <correlation name>s or <table or query name>s of the <table reference> bound to that row. The <search condition>s are effectively evaluated for *RI* before updating any row of *T* and prior to the invocation of any <triggered action> caused by the update of any row of *T* and before inserting any rows into *T* and prior to the invocation of any <triggered action> caused by the insert of any row of *T*.

Case:

- A) If <target table> contains ONLY, then *RI* is a subject row if *RI* has no subrow in a proper subtable of *T* and the result of the <search condition>s are True for some row *R2* of *Q* and *RI* is not a subject row identified by any other <merge when matched

clause> that precedes *MWMC* in the <merge operation specification>. *R2* is the matching row.

- B) Otherwise, *R1* is a subject row if the result of the <search condition>s are *True* for some row *R2* of *Q* and *R1* is not a subject row identified by any other <merge when matched clause> that precedes *MWMC* in the <merge operation specification>. *R2* is the matching row.

2) If *R1* is a subject row, then:

- A) Let *M* be the number of matching rows in *Q* for *R1*.
- B) If *M* is greater than 1 (one), then an exception condition is raised: *cardinality violation*.
- C) The <update source> of each <set clause> is effectively evaluated for *R1* before any row of *T* is updated and prior to the invocation of any <triggered action> caused by the update of any row of *T*. The resulting value is the update value.
- D) A candidate new row is constructed by copying the subject row and updating it as specified by each <set clause> by applying the General Rules of Subclause 14.15, “<set clause list>”.

- ii) If *T* is a base table, then each subject row is also an object row; otherwise, an object row is any row of a leaf generally underlying table of *T* from which a subject row is derived.

NOTE 395 — The data values allowable in the object rows may be constrained by a WITH CHECK OPTION constraint. The effect of a WITH CHECK OPTION constraint is defined in the General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”.

- iii) If any row in the set of object rows has been marked for deletion by any <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned> that identifies some open cursor *CR* or updated by any <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies some open cursor *CR*, then a completion condition is raised: *warning — cursor operation conflict*.

- iv) Let *CL* be the columns of *T* identified by the <object column>s contained in the <set clause list>.

- v) Each subject row *SR* is identified for replacement, by its corresponding candidate new row *CNR*, in *T*. The set of (*SR*, *CNR*) pairs is the replacement set for *T*.

NOTE 396 — Identifying a row for replacement, associating a replacement row with an identified row, and associating a replacement set with a table are implementation-dependent operations.

- vi) Case:

1) If *T* is a base table, then:

A) Case:

- I) If <target table> specifies ONLY, then *T* is *identified for replacement processing without subtables* with respect to object columns *CL*.
- II) Otherwise, *T* is *identified for replacement processing with subtables* with respect to object columns *CL*.

NOTE 397 — Identifying a base table for replacement processing, with or without subtables, is an implementation-dependent mechanism. In general, though not here, the list of object columns can be empty.

- B) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.
- 2) If *T* is a viewed table, then the General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”, are applied with <target table> as *VIEW NAME* and the replacement set for *T* as *REPLACEMENT SET FOR VIEW NAME*.
- b) If <merge when not matched clause> *MWNMC* is specified, then:
- i) Let *TR1* be the <target table> immediately contained in <merge statement>, let *TR2* be the <table reference> immediately contained in <merge statement>, and let *SC1* be the <search condition> immediately contained in <merge statement>. If <merge correlation name> is specified, then let *MCN* be “AS <merge correlation name>”; otherwise, let *MCN* be a zero-length string. If *MWNMC* immediately contains a <search condition> *SC2*, then let *ONSC2* be “OR NOT *SC2*”; otherwise, let *ONSC2* be a zero-length string. Let *S1* be the result of
 

```
SELECT *
FROM TR1 MCN, TR2
WHERE SC1 ONSC2
```
  - ii) Let *S2* be the collection of rows of *Q* for which there exists in *S1* some row that is the concatenation of some row *R1* of *T* and some row *R2* of *Q*.
  - iii) Let *S3* be the collection of rows of *Q* that are not in *S2*. Let *SN3* be the effective distinct name for *S3*. Let *EN* be the exposed <correlation name> or <table or query name> of *TR2*.
  - iv) Let *S4* be the result of:
 

```
SELECT EXP1, EXP2, . . . , EXPNI
FROM SN3 AS EN
```
  - v) Let *S5* be the collection of rows of *S4* for which no candidate rows have been effectively created by any other <merge when not matched clause> that precedes *MWNMC* in the <merge operation specification>.
  - vi) *S5* is effectively evaluated before insertion of any rows into or update of any rows in *T*.
  - vii) For each row *R* of *S5*:
    - 1) A candidate row of *T* is effectively created in which the value of each column is its default value, as specified in the General Rules of Subclause 11.5, “<default clause>”. The candidate row consists of every column of *T*.
    - 2) If *T* has a column *RC* of which some underlying column is a self-referencing column, then
 

Case:

      - A) If *RC* is a system-generated self-referencing column, then the value of *RC* is effectively replaced by the REF value of the candidate row.
      - B) If *RC* is a derived self-referencing column, then the value of *RC* is effectively replaced by a value derived from the columns in the candidate row that correspond to the list of

attributes of the derived representation of the reference type of  $RC$  in an implementation-dependent manner.

- 3) For each object column in the candidate row, let  $C_i$  be the object column identified by the  $i$ -th <column name> in the <insert column list> and let  $SV_i$  be the  $i$ -th value of  $R$ .
- 4) For every  $C_i$  for which one of the following conditions is true:
  - A)  $C_i$  is not marked as unassigned and no underlying column of  $C_i$  is a self-referencing column.
  - B) Some underlying column of  $C_i$  is a user-generated self-referencing column.
  - C) Some underlying column of  $C_i$  is a self-referencing column and OVERRIDING SYSTEM VALUE is specified.
  - D) Some underlying column of  $C_i$  is an identity column and the  $i$ -th column of  $R$  is not derived from <default specification> and OVERRIDING SYSTEM VALUE is specified.
  - E) Some underlying column of  $C_i$  is an identity column whose descriptor includes an indication that values are generated by default and neither OVERRIDING USER VALUE is specified nor is the  $i$ -th column derived from <default specification>.

the General Rules of Subclause 9.2, “Store assignment”, are applied to  $C_i$  and  $SV_i$  as *TARGET* and *SOURCE*, respectively.  $C_i$  is no longer marked as unassigned.

NOTE 398 — The data values allowable in the candidate row may be constrained by a WITH CHECK OPTION constraint. The effect of a WITH CHECK OPTION constraint is defined in the General Rules of Subclause 15.12, “Effect of inserting a table into a viewed table”.

viii) Let  $S$  be the table consisting of the candidate rows.

Case:

- 1) If  $T$  is a base table, then:
  - A)  $T$  is *identified for insertion of source table*  $S$ .

NOTE 399 — Identifying a base table for insertion of a source table is an implementation-dependent operation.
  - B) The General Rules of Subclause 15.10, “Effect of inserting tables into base tables”, are applied.
- 2) If  $T$  is a viewed table, then the General Rules of Subclause 15.12, “Effect of inserting a table into a viewed table”, are applied with  $S$  as *SOURCE* and  $T$  as *TARGET*.

7) If  $Q$  is empty, then a completion condition is raised: *no data*.

## Conformance Rules

- 1) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <merge statement> in which a leaf generally underlying table of  $T$  is generally contained in a <query expression> immediately contained in the <table reference> except as the <table or query name> or <correlation name> of a column reference.

- 2) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <merge statement> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in the <search condition>.
- 3) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <merge statement> that does not satisfy the condition: for each column *C* identified in the explicit or implicit <insert column list>, if the declared type of *C* is a structured type *TY*, then the declared type of the corresponding column of the <query expression> or <contextually typed table value constructor> is *TY*.
- 4) Without Feature F312, “MERGE statement”, conforming SQL language shall not contain a <merge statement>.
- 5) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain a <merge statement> that contains an <target table> that identifies a table that is not simply updatable.
- 6) Without Feature F313, “Enhanced MERGE statement”, in conforming SQL language, a <merge statement> shall not contain each of <merge when matched clause> and <merge when not matched clause> more than once.
- 7) Without Feature F313, “Enhanced MERGE statement”, in conforming SQL language, a <merge when matched clause> or a <merge when not matched clause> shall not immediately contain a <search condition>.

## 14.13 <update statement: positioned>

*This Subclause is modified by Subclause 12.7, “<update statement: positioned>”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 11.13, “<update statement: positioned>”, in ISO/IEC 9075-10.*

*This Subclause is modified by Subclause 14.6, “<update statement: positioned>”, in ISO/IEC 9075-14.*

### Function

Update a row of a table.

### Format

```

10 14 <update statement: positioned> ::=
 UPDATE <target table> [[AS] <correlation name>]
 SET <set clause list>
 WHERE CURRENT OF <cursor name>

```

### Syntax Rules

- 1) 04 Let *USP* be the <update statement: positioned> and let *CN* be the <cursor name> immediately contained in *USP*. *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*.
- 2) *CN* shall identify a standing cursor.
- 3) Let *CDD* be the cursor declaration descriptor of the standing cursor identified by *CN*.
- 4) The cursor specification of *CDD* shall be updatable.
- 5) Let *TU* be the simply underlying table of *CDD*. *TU* is the *subject table* of the <update statement: positioned>. Let *LUT* be the leaf underlying table of *TU* such that *TU* is one-to-one with respect to *LUT*.
- 6) Let *TN* be the <table name> contained in <target table> *TT*. *TN* shall identify *LUT*.
- 7) <target table> shall specify ONLY if and only if the <table reference> contained in *TU* that references *LUT* specifies ONLY.
- 8) *TN* shall not identify an old transition table or a new transition table.
- 9) Let *T* be the table identified by *TN*.
- 10) Case:
  - a) If <correlation name> is specified, then let *CN* be that <correlation name>.
  - b) Otherwise, let *CN* be the <table name> contained in <target table>. *CN* is an exposed <table or query name>.
- 11) The scope of *CN* is <set clause list>.
- 12) If the declared <cursor specification> of *CDD* is ordered, then for each <object column> *OC* contained in <set clause list>, the <order by clause> simply contained in the <cursor specification> of *CDD* shall not generally contain a <column reference> that references *OC* or an underlying column of the column identified by *OC*.



**14.13 <update statement: positioned>**

- 13) Each <column name> specified as an <object column> shall identify a column in the explicit or implicit <column name list> contained in the explicit or implicit <updatability clause> of the <cursor specification> of *CDD*.

**Access Rules**

- 1) Case:
  - a) If <update statement: positioned> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema. The applicable privileges for *A* shall include UPDATE for each <object column>.
  - b) Otherwise, the current privileges shall include UPDATE for each <object column>.

**General Rules**

- 1) Let *CR* be the cursor instance descriptor of the current SQL-session whose cursor declaration descriptor is *CDD*.
- 2) Let *SCL* be the <set clause list>.
- 3) The General Rules of Subclause 15.6, “Effect of a positioned update”, are applied with *CR* as *CURSOR*, *SCL* as *SET CLAUSE LIST*, *USP* as *STATEMENT*, and *TT* as *TARGET*.

**Conformance Rules**

- 1) Without Feature F831, “Full cursor update”, conforming SQL language shall not contain an <update statement: positioned> in which the declared <cursor specification> of *CDD* is ordered.

## 14.14 <update statement: searched>

This Subclause is modified by *Subclause 14.7*, “<update statement: searched>”, in ISO/IEC 9075-14.

### Function

Update rows of a table.

### Format

```
14 <update statement: searched> ::=
 UPDATE <target table> [[AS] <correlation name>]
 SET <set clause list>
 [WHERE <search condition>]
```

### Syntax Rules

- 1) Let *TN* be the <table name> contained in <target table>; let *T* be the table identified by *TN*. *T* shall be an updatable table or a trigger-updatable table.
- 2) *T* is the *subject table* of the <update statement: searched>.
- 3) *TN* shall not identify an old transition table or a new transition table.
- 4) Case:
  - a) If <correlation name> is specified, then let *CN* be that <correlation name>.
  - b) Otherwise, let *CN* be the <table name> contained in <target table>. *CN* is an exposed <table or query name>.
- 5) The scope of *CN* is <set clause list> and <search condition>.
- 6) If the <update statement: searched> is contained in a <triggered SQL statement>, then the <search condition> shall not contain a <value specification> that specifies a parameter reference.
- 7) The <search condition> shall not generally contain a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.

### Access Rules

- 1) Case:
  - a) If <update statement: searched> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema.
    - i) The applicable privileges for *A* for *TN* shall include UPDATE for each <object column>.
    - ii) If <target table> immediately contains ONLY, then the applicable privileges for *A* shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *T*.

- b) Otherwise,
  - i) The current privileges for *TN* shall include UPDATE for each <object column>.
  - ii) If <target table> immediately contains ONLY, then the current privileges shall include SELECT WITH HIERARCHY OPTION on at least one supertable of *T*.

## General Rules

- 1) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only and *T* is not a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 2) If there is any sensitive cursor *CR* that is currently open in the SQL-transaction in which this SQL-statement is being executed, then

Case:

- a) If *CR* has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of this statement shall be made visible to *CR* or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - b) Otherwise, whether the change resulting from the successful execution of this SQL-statement is made visible to *CR* is implementation-defined.
- 3) If there is any open, insensitive cursor *CR*, then either the change resulting from the successful execution of this statement shall be invisible to *CR*, or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - 4) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.

- 5) Case:

- a) If <target table> contains ONLY, then

Case:

- i) If a <search condition> is not specified, then all rows of *T* for which there is no subrow in a proper subtable of *T* are the subject rows.
- ii) If a <search condition> is specified, then it is effectively evaluated for each row of *T* with the exposed <correlation name>s or <table or query name>s of the <table reference> bound to that row, and the subject rows are those rows for which the result of the <search condition> is True and for which there is no subrow in a proper subtable of *T*. The <search condition> is effectively evaluated for each row of *T* before updating any row of *T*.

- b) Otherwise,

Case:

- i) If a <search condition> is not specified, then all rows of *T* are the subject rows.
- ii) If a <search condition> is specified, then it is effectively evaluated for each row of *T* with the exposed <table name> of the <target table> bound to that row, and the subject rows are those

## 14.14 &lt;update statement: searched&gt;

rows for which the result of the <search condition> is *True*. The <search condition> is effectively evaluated for each row of *T* before any row of *T* is updated.

- 6) If *T* is a base table, then each subject row is also an *object row*; otherwise, an *object row* is any row of a leaf generally underlying table of *T* from which a subject row is derived.
- 7) If any row in the set of object rows has been marked for deletion by any <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned> that identifies some open cursor *CR* or updated by any <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies some open cursor *CR*, then a completion condition is raised: *warning — cursor operation conflict*.
- 8) If a <search condition> is specified, then the <search condition> is evaluated for each row of *T* prior to the invocation of any <triggered action> caused by the update of any row of *T*.
- 9) The <update source> of each <set clause> is effectively evaluated for each row of *T* before any row of *T* is updated.
- 10) For each subject row, a candidate new row is constructed by copying the subject row and updating it as specified by each <set clause> by applying the General Rules of Subclause 14.15, “<set clause list>”.

NOTE 400 — The data values allowable in the object rows may be constrained by a WITH CHECK OPTION constraint. The effect of a WITH CHECK OPTION constraint is defined in the General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”.

- 11) Let *CL* be the columns of *T* identified by the <object column>s contained in the <set clause list>.
- 12) Each subject row *SR* is identified for replacement, by its corresponding candidate new row *CNR*, in *T*. The set of (*SR*, *CNR*) pairs is the *replacement set* for *T*.

NOTE 401 — Identifying a row for replacement, associating a replacement row with an identified row, and associating a replacement set with a table are implementation-dependent operations.

## 13) Case:

- a) If *T* is a base table, then:

- i) Case:

- 1) If <target table> specifies ONLY, then *T* is *identified for replacement processing without subtables* with respect to object columns *CL*.
- 2) Otherwise, *T* is *identified for replacement processing with subtables* with respect to object columns *CL*.

NOTE 402 — Identifying a base table for replacement processing, with or without subtables, is an implementation-dependent mechanism. In general, though not here, the list of object columns can be empty.

- ii) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.

- b) If *T* is a viewed table, then the General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”, are applied with <target table> as *VIEW NAME* and the replacement set for *T* as *REPLACEMENT SET FOR VIEW NAME*.

- 14) If the set of object rows is empty, then a completion condition is raised: *no data*.

## Conformance Rules

- 1) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain an <update statement: searched> in which a leaf generally underlying table of  $T$  is an underlying table of any <query expression> generally contained in the <search condition>.
- 2) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain an <update statement: searched> that contains a <target table> that identifies a table that is not simply updatable.

## 14.15 <set clause list>

### Function

Specify a list of updates.

### Format

```

<set clause list> ::=
 <set clause> [{ <comma> <set clause> }...]

<set clause> ::=
 <multiple column assignment>
 | <set target> <equals operator> <update source>

<set target> ::=
 <update target>
 | <mutated set clause>

<multiple column assignment> ::=
 <set target list> <equals operator> <assigned row>

<set target list> ::=
 <left paren> <set target> [{ <comma> <set target> }...] <right paren>

<assigned row> ::=
 <contextually typed row value expression>

<update target> ::=
 <object column>
 | <object column>
 <left bracket or trigraph> <simple value specification> <right bracket or trigraph>

<object column> ::=
 <column name>

<mutated set clause> ::=
 <mutated target> <period> <method name>

<mutated target> ::=
 <object column>
 | <mutated set clause>

<update source> ::=
 <value expression>
 | <contextually typed value specification>

```

### Syntax Rules

- 1) Let *T* be the table identified by the <target table> contained in the containing <update statement: positioned>, <update statement: searched>, or <merge statement>.
- 2) If *T* is not trigger updatable, then each <column name> specified as an <object column> shall identify an updatable column of *T*.

NOTE 403 — The notion of updatable columns of base tables is defined in Subclause 4.14, “Tables”. The notion of updatable columns of viewed tables is defined in Subclause 11.23, “<view definition>”.

- 3) Each <set clause> *SC* that immediately contains a <multiple column assignment> is effectively replaced by a <set clause list> *MSCL* as follows:
- Let *STN* be the number of <set target>s contained in <set target list>.
  - STN* shall be equal to the degree of the <assigned row> *AR* contained in *SC*.
  - Let *ST<sub>i</sub>*, 1 (one) ≤ *i* ≤ *STN*, be the *i*-th <set target> contained in the <set target list> of *SC* and let *DT<sub>i</sub>* be the declared type of the *i*-th field of *AR*. The *i*-th <set clause> in *MSCL* is:

$$ST_i =$$

```
CAST (AR AS ROW (F1 DT1,
F2 DT2, . . . ,
FSTN DTSTN)) . Fi
```

NOTE 404 — “Fn” here stands for the <field name> consisting of the letter “F” followed, with no intervening <separator> by the decimal <digit> or <digit>s comprising a <literal> corresponding to the value *n*.

- If <set clause> *SC* specifies an <object column> that references a column of which some underlying column is either a generated column or an identity column whose descriptor indicates that values are always generated, then the <update source> specified in *SC* shall consist of a <default specification>.
- A <value expression> simply contained in an <update source> in a <set clause> shall not directly contain a <set function specification>.
- If the <set clause list> *OSCL* contains one or more <set clause>s that contain a <mutated set clause>, then:
  - Let *N* be the number of <set clause>s in *OSCL* that contain a <mutated set clause>.
  - For 1 (one) ≤ *i* ≤ *N*:
    - Let *SC<sub>i</sub>* be the *i*-th <set clause> that contains a <mutated set clause>.
    - Let *RCVE<sub>i</sub>* be the <update source> immediately contained in *SC<sub>i</sub>*.
    - Let *MSC<sub>i</sub>* be the <mutated set clause> immediately contained in the <set target> immediately contained in *SC<sub>i</sub>*.
    - Let *OC<sub>i</sub>* be the <object column> contained in *MSC<sub>i</sub>*. The declared type of the column identified by *OC<sub>i</sub>* shall be a structured type.
    - Let *M<sub>i</sub>* be the number of <method name>s contained in *MSC<sub>i</sub>*.
    - For 1 (one) ≤ *j* ≤ *M<sub>i</sub>*:
      - If *j* = 1 (one), then

Case:

        - Let *MT<sub>i,1</sub>* be the <mutated target> immediately contained in *MSC<sub>i</sub>*.
        - Let *MN<sub>i,1</sub>* be the <method name> immediately contained in *MSC<sub>i</sub>*.
        - Let *V<sub>i,1</sub>* be:

$$MT_{i,1} \cdot MN_{i,1} (RCVE_i)$$

2) Otherwise:

A) Let  $MT_{i,j}$  be the <mutated target> immediately contained in the <mutated set clause> immediately contained in  $MT_{i,j-1}$ .

B) Let  $MN_{i,j}$  be the <method name> immediately contained in the <mutated set clause> immediately contained in  $MT_{i,j-1}$ .

C) Let  $V_{i,j}$  be

$$MT_{i,j} \cdot MN_{i,j} (V_{i,j-1})$$

c) *OSCL* is equivalent to a <set clause list> *NSCL* derived as follows:

i) Let *NSCL* be a <set clause list> derived from *OSCL* by replacing every <set clause>  $SC_a$ , 1 (one)  $\leq a \leq N$ , that contains a <mutated set clause> with:

$$MT_{a,Ma} = V_{a,Ma}$$

ii) For 1 (one)  $\leq b \leq N$ , if there exists a  $c$  such that  $c < b$  and  $OC_c$  is equivalent to  $OC_b$ , then:

1) Every occurrence of  $OC_b$  in  $V_{b,Mb}$  is replaced by  $V_{c,Mc}$ .

2)  $SC_c$  is deleted from *NSCL*.

7) Equivalent <object column>s shall not appear more than once in a <set clause list>.

NOTE 405 — Multiple occurrences of equivalent <object column>s within <mutated set clause>s are eliminated by the preceding Syntax Rule of this Subclause.

8) If the <update source> of <set clause>  $SC$  specifies a <contextually typed value specification> *CVS*, then the data type of *CVS* is the data type *DT* of the <update target> or <mutated set clause> specified in  $SC$ .

9) If *CVS* is an <empty specification>, then *DT* shall be a collection type. If *CVS* specifies ARRAY, then *DT* shall be an array type. If *CVS* specifies MULTISSET, then *DT* shall be a multiset type.

10) For every <object column> in a <set clause>,

Case:

- a) If the <update target> immediately contains <simple value specification>, then the declared type of the column of *T* identified by the <object column> shall be an array type. The Syntax Rules of Subclause 9.2, “Store assignment”, apply to an arbitrary site whose declared type is the element type of the column of *T* identified by the <object column> and the <update source> of the <set clause> as *TARGET* and *VALUE*, respectively.
- b) Otherwise, the Syntax Rules of Subclause 9.2, “Store assignment”, apply to the column of *T* identified by the <object column> and the <update source> of the <set clause> as *TARGET* and *VALUE*, respectively.



## Access Rules

*None.*

## General Rules

- 1) A <set clause> specifies one or more object columns and an update value. An *object column* is a column identified by an <object column> in the <set clause>. The *update value* is the value specified by the <update source> contained in the <set clause>.

- 2) The value of the *i*-th object column denoted by *C*, is replaced as follows.

Case:

- a) If the *i*-th <set clause> contains an <update target> that immediately contains a <simple value specification>, then

Case:

- i) If the value of *C* is the null value, then an exception condition is raised: *data exception — null value in array target*.

- ii) Otherwise:

- 1) Let *N* be the maximum cardinality of *C*.

- 2) Let *M* be the cardinality of the value of *C*.

- 3) Let *I* be the value of the <simple value specification> immediately contained in <update target>.

- 4) Let *EDT* be the element type of *C*.

- 5) Case:

- A) If *I* is greater than zero and less than or equal to *M*, then the value of *C* is replaced by an array *A* with element type *EDT* and cardinality *M* derived as follows:

- I) For *j* varying from 1 (one) to *I*–1 and from *I*+1 to *M*, the *j*-th element in *A* is the value of the *j*-th element in *C*.

- II) The *I*-th element of *A* is set to the *i*-th update value, denoted by *SV*, by applying the General Rules of [Subclause 9.2, “Store assignment”](#), to the *I*-th element of *A* and *SV* as *TARGET* and *VALUE*, respectively.

- B) If *I* is greater than *M* and less than or equal to *N*, then the value of *C* is replaced by an array *A* with element type *EDT* and cardinality *I* derived as follows:

- I) For *j* varying from 1 (one) to *M*, the *j*-th element in *A* is the value of the *j*-th element in *C*.

- II) For *j* varying from *M*+1 to *I*–1, the *j*-th element in *A* is the null value.

- III) The *I*-th element of *A* is set to the *i*-th update value, denoted by *SV*, by applying the General Rules of [Subclause 9.2, “Store assignment”](#), to the *I*-th element of *A* and *SV* as *TARGET* and *VALUE*, respectively.

- C) Otherwise, an exception condition is raised: *data exception — array element error*.
- b) Otherwise, the value of *C* is replaced by the *i*-th update value, denoted by *SV*. The General Rules of Subclause 9.2, “Store assignment”, are applied to *C* and *SV* as *TARGET* and *VALUE*, respectively.

## Conformance Rules

- 1) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <set clause> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in any <value expression> simply contained in an <update source> or <assigned row> immediately contained in the <set clause>.
- 2) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <update target> that immediately contains a <simple value specification>.
- 3) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <set clause> in which the declared type of the <update target> in the <set clause> is a structured type *TY* and the declared type of the <update source> or corresponding field of the <assigned row> contained in the <set clause> is not *TY*.
- 4) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <set clause> that contains a <mutated set clause> and in which the declared type of the last <method name> identifies a structured type *TY*, and the declared type of the <update source> contained in the <set clause> is not *TY*.
- 5) Without Feature T641, “Multiple column assignment”, conforming SQL language shall not contain a <multiple column assignment>.

## 14.16 <temporary table declaration>

This Subclause is modified by *Subclause 12.8*, “<temporary table declaration>”, in ISO/IEC 9075-4.

### Function

Declare a declared local temporary table.

### Format

```
<temporary table declaration> ::=
 DECLARE LOCAL TEMPORARY TABLE <table name> <table element list>
 [ON COMMIT <table commit action> ROWS]
```

### Syntax Rules

- 1) Let *TN* be the <table name> of a <temporary table declaration> *TTD*, and let *T* be the <qualified identifier> of *TN*.
- 2) 04 *TTD* shall be contained in an <SQL-client module definition>.
- 3) Case:
  - a) If *TN* contains a <local or schema qualifier> *LSQ*, then *LSQ* shall be “MODULE”.
  - b) If *TN* does not contain a <local or schema qualifier>, then “MODULE” is implicit.
- 4) 04 If a <temporary table declaration> is contained in an <SQL-client module definition> *M*, then the <qualified identifier> of *TN* shall not be equivalent to the <qualified identifier> of the <table name> of any other <temporary table declaration> that is contained in *M*.
- 5) The descriptor of the table defined by a <temporary table declaration> includes *TN* and the column descriptor specified by each <column definition>. The *i*-th column descriptor is given by the *i*-th <column definition>.
- 6) <table element list> shall contain at least one <column definition> or at least one <like clause>.
- 7) If ON COMMIT is not specified, then ON COMMIT DELETE ROWS is implicit.

### Access Rules

*None.*

### General Rules

- 1) 04 Let *U* be the implementation-dependent <schema name> that is effectively derived from the implementation-dependent SQL-session identifier associated with the SQL-session and an implementation-dependent name associated with the SQL-client module that contains the <temporary table declaration>.
- 2) Let *UI* be the current user identifier and let *R* be the current role name.

**14.16 <temporary table declaration>**

Case:

- a) If *UI* is not the null value, then let *A* be *UI*.
  - b) Otherwise, let *A* be *R*.
- 3) 04 The definition of *T* within an SQL-client module is effectively equivalent to the definition of a persistent base table *U.T*. Within the SQL-client module, any reference to *MODULE.T* is equivalent to a reference to *U.T*.
  - 4) A set of privilege descriptors is created that define the privileges INSERT, SELECT, UPDATE, DELETE, and REFERENCES on this table and INSERT, SELECT, UPDATE, and REFERENCES for every <column definition> in the table definition to *A*. These privileges are not grantable. The grantor for each of these privilege descriptors is set to the special grantor value “\_SYSTEM”. The grantee is “PUBLIC”.
  - 5) The definition of a temporary table persists for the duration of the SQL-session. The termination of the SQL-session is effectively followed by the execution of the following <drop table statement> with the current authorization identifier *A* and current <schema name> *U* without further Access Rule checking:

```
DROP TABLE T CASCADE
```

- 6) The definition of a declared local temporary table does not appear in any view of the Information Schema.

NOTE 406 — The Information Schema is defined in [ISO9075-11].

## Conformance Rules

- 1) Without Feature F531, “Temporary tables”, conforming SQL language shall not contain a <temporary table declaration>.

## 14.17 <free locator statement>

### Function

Remove the association between a locator variable and the value that is represented by that locator.

### Format

```
<free locator statement> ::=
 FREE LOCATOR <locator reference> [{ <comma> <locator reference> }...]

<locator reference> ::=
 <host parameter name>
 | <embedded variable name>
 | <dynamic parameter specification>
```

### Syntax Rules

- 1) Each host parameter identified by <host parameter name> immediately contained in <locator reference> shall be a binary large object locator parameter, a character large object locator parameter, an array locator parameter, a multiset locator parameter, or a user-defined type locator parameter.
- 2) Each host variable identified by the <embedded variable name> immediately contained in <locator reference> shall be a binary object locator variable, a character large object locator variable, an array locator variable, a multiset locator parameter, or a user-defined type locator variable.

### Access Rules

*None.*

### General Rules

- 1) For every <locator reference> *LR* immediately contained in <free locator statement>, let *L* be the value of *LR*.  
Case:
  - a) If *L* is not a valid locator value, then an exception condition is raised: *locator exception — invalid specification*.
  - b) Otherwise, *L* is marked invalid.

### Conformance Rules

- 1) Without Feature T561, “Holdable locators”, conforming SQL language shall not contain a <free locator statement>.

## 14.18 <hold locator statement>

### Function

Mark a locator variable as being holdable.

### Format

```
<hold locator statement> ::=
 HOLD LOCATOR <locator reference> [{ <comma> <locator reference> }...]
```

### Syntax Rules

- 1) Each host parameter identified by <host parameter name> immediately contained in <locator reference> shall be a binary large object locator parameter, a character large object locator parameter, an array locator parameter, a multiset locator parameter, or a user-defined type locator parameter.

### Access Rules

*None.*

### General Rules

- 1) For every <locator reference> *LR* immediately contained in <hold locator statement>, let *L* be the value of *LR*.

Case:

- a) If *L* is not a valid locator value, then an exception condition is raised: *locator exception — invalid specification*.
- b) Otherwise, *L* is marked *holdable*.

### Conformance Rules

- 1) Without Feature T561, “Holdable locators”, conforming SQL language shall not contain a <hold locator statement>.

## 15 Additional data manipulation rules

### 15.1 Effect of opening a cursor

*This Subclause is modified by Subclause 7.1, “Effect of opening a cursor”, in ISO/IEC 9075-3.  
This Subclause is modified by Subclause 13.1, “Effect of opening a cursor”, in ISO/IEC 9075-4.*

#### Function

Specify the effect of opening a cursor that is not a received cursor.

#### Syntax Rules

*None.*

#### Access Rules

*None.*

#### General Rules

- 1) Let *CR* be the *CURSOR* specified in an application of this Subclause. If *CR* is not in the closed state, then an exception condition is raised: *invalid cursor state*.
- 2) Let *CDD* be the cursor declaration descriptor of *CR*.
- 3) Case:
  - a) 03 If the kind of cursor described by *CDD* is a standing cursor, then let *S* be the declared <cursor specification> of *CDD*.
  - b) If the kind of cursor described by *CDD* is a declared dynamic cursor, then let *S* be the prepared statement indicated by the <statement name> that is the origin of *CDD*.
  - c) If the kind of cursor described by *CDD* is an extended statement cursor, then let *S* be the prepared statement that is the origin of *CDD*.

NOTE 407 — If the kind of cursor is a received cursor, then this Subclause does not apply; instead, Subclause 15.2, “Effect of receiving a result set”, applies.
- 4) *CR* is opened, and a result set descriptor *RSD* is created and included in *CR*, in the following steps:
  - a) The <cursor specification> of *RSD* is a copy *CS* of *S* that is effectively created as follows:
    - i) 04 Each <embedded variable specification>, <host parameter specification>, <SQL parameter reference>, and <dynamic parameter specification> is replaced by a <literal> denoting the value

## 15.1 Effect of opening a cursor

resulting from evaluating the <embedded variable specification>, <host parameter specification>, <SQL parameter reference>, and <dynamic parameter specification>, respectively, with all such evaluations effectively done at the same instant in time.

- ii) Each <value specification> generally contained in *S* that is `CURRENT_USER`, `CURRENT_ROLE`, `SESSION_USER`, `SYSTEM_USER`, `CURRENT_CATALOG`, `CURRENT_SCHEMA`, `CURRENT_PATH`, `CURRENT_DEFAULT_TRANSFORM_GROUP`, or `CURRENT_TRANSFORM_GROUP_FOR_TYPE` <path-resolved user-defined type name> is replaced by a <literal> denoting the value resulting from evaluation of `CURRENT_USER`, `CURRENT_ROLE`, `SESSION_USER`, `SYSTEM_USER`, `CURRENT_CATALOG`, `CURRENT_SCHEMA`, `CURRENT_PATH`, `CURRENT_DEFAULT_TRANSFORM_GROUP`, or `CURRENT_TRANSFORM_GROUP_FOR_TYPE` <path-resolved user-defined type name>, respectively, with all such evaluations effectively done at the same instant in time.
  - iii) Each <datetime value function> generally contained in *S* is replaced by a <literal> denoting the value resulting from evaluation of that <datetime value function>, with all such evaluations effectively done at the same instant in time.
- b) Case:
- i) If *CR* is a standing cursor, then the operational properties of *RSD* are the same as the corresponding declared properties of *CDD*.
  - ii) Otherwise, *CR* is a dynamic cursor.
    - 1) The operational properties of *RSD* are initially copied from the corresponding declared properties of *CDD*.
    - 2) If the <prepare statement> that prepared *S* contained an <attributes variable>, then let *CA* be the value of that <attributes variable>.
      - A) If *CA* contains a <cursor sensitivity>, then the operational sensitivity property of *RSD* is set to that <cursor sensitivity>.
      - B) If *CA* contains a <cursor scrollability>, then the operational scrollability property of *RSD* is set to that <cursor scrollability>.
      - C) If *CA* contains a <cursor holdability>, then the operational holdability property of *RSD* is set to that <cursor holdability>.
      - D) If *CA* contains a <cursor returnability>, then the operational returnability property of *RSD* is set to that <cursor returnability>.
- c) If *CR* is being opened during either
- i) the execution of an SQL-invoked function, or
  - ii) the execution of an externally-invoked procedure
- without the intervening execution of an SQL-invoked procedure, then the operational returnability property of *RSD* is set to `WITHOUT RETURN`.
- d) Let *QE* be the <query expression> simply contained in *CS*.
  - e) Let *T* be the table specified by *QE*.
  - f) Case:



- i) If the operational sensitivity property of *RSD* is INSENSITIVE, then let *TT* be a copy of *T*.
  - ii) Otherwise, let *TT* be *T*.
  - g) The sequence of rows of *RSD* consists of the rows of *TT*, ordered as determined by the General Rules of Subclause 7.13, “<query expression>”.
- 5) If *CR* is insensitive, and the SQL-implementation is unable to guarantee that significant changes will be invisible through *CR* during the SQL-transaction in which *CR* is opened and every subsequent SQL-transaction during which it may be held open, then an exception condition is raised: *cursor sensitivity exception — request rejected*.
- 6) If *CR* is sensitive, and the SQL-implementation is unable to guarantee that significant changes will be visible through *CR* during the SQL-transaction in which *CR* is opened, then an exception condition is raised: *cursor sensitivity exception — request rejected*.
- NOTE 408 — The visibility of significant changes through a sensitive holdable cursor during a subsequent SQL-transaction is implementation-defined.
- 7) Whether an SQL-implementation is able to disallow significant changes that would not be visible through a currently open cursor is implementation-defined.
- 8) If the operational returnability property of *RSD* is WITH RETURN, then let *SIP* be the active SQL-invoked procedure, let *INV* be the invoker of *SIP*, and let *RSS* be the result set sequence for *SIP* and *INV* in the active SQL-session context. *RSD* is added to the end of *RSS*.

## Conformance Rules

*None.*

## 15.2 Effect of receiving a result set

### Function

Advance a received cursor to the next result set in a result set sequence.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *CR* be the *CURSOR* and *RSS* be the *RESULT SET SEQUENCE* in an application of this Subclause.
- 2) Let *RS* be the first result set descriptor in *RSS*.
- 3) *CR* is placed in the open state, with *RS* as the result set descriptor included in *CR*.
- 4) *RS* is removed from *RSS*.
- 5) If the operational scrollability property of *RS* is NO SCROLL, then the position of *RS* is set to before the first row.
- 6) The operational returnability property of *RS* is implementation-defined.
- 7) The <cursor specification> of *RS* is made not updatable by replacing the explicit or implicit <updatability clause> with FOR READ ONLY.

### Conformance Rules

*None.*

## 15.3 Determination of the current row of a cursor

### Function

Specify how the current row of a cursor is determined.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *CR* be the *CURSOR* and let *FO* be the *FETCH ORIENTATION* specified in an application of this Subclause.
- 2) If *CR* is not in the open state, then an exception condition is raised: *invalid cursor state*.
- 3) Let *RSD* be the result set descriptor of *CR*.
- 4) If *FO* is not NEXT and the operational scrollability property of *RSD* is not SCROLL, then an exception condition is raised: *syntax error or access rule violation*.
- 5) Case:
  - a) If *FO* contains a <simple value specification>, then let *J* be the value of that <simple value specification>.
  - b) If *FO* specifies NEXT or FIRST, then let *J* be +1.
  - c) If *FO* specifies PRIOR or LAST, then let *J* be -1.
- 6) Let *T* be the sequence of rows of *RSD*.
- 7) Let *T<sub>f</sub>* be a sequence of rows of the same degree as *T*.
 

Case:

  - a) If *FO* specifies ABSOLUTE, FIRST, or LAST, then let *T<sub>f</sub>* contain all rows of *T*, preserving their order in *T*.
  - b) If *FO* specifies NEXT or specifies RELATIVE with a positive value of *J*, then:
    - i) If *T* is empty or if the position of *CR* is on or after the last row of *T*, then let *T<sub>f</sub>* be a sequence of no rows.
    - ii) If the position of *CR* is on a row *R* that is other than the last row of *T*, then let *T<sub>f</sub>* contain all rows of *T* ordered after row *R*, preserving their order in *T*.

### 15.3 Determination of the current row of a cursor

- iii) If the position of *CR* is before a row *R*, then let  $T_t$  contain row *R* and all rows of *T* ordered after row *R*, preserving their order in *T*.
  - c) If *FO* specifies PRIOR or specifies RELATIVE with a negative value of *J*, then:
    - i) If *T* is empty or if the position of *CR* is on or before the first row of *T*, then let  $T_t$  be a sequence of no rows.
    - ii) If the position of *CR* is on a row *R* that is other than the first row of *T*, then let  $T_t$  contain all rows of *T* ordered before row *R*, preserving their order in *T*.
    - iii) If the position of *CR* is before a row *R* that is not the first row of *T*, then let  $T_t$  contain row all rows of *T* ordered before row *R*, preserving their order in *T*.
    - iv) If the position of *CR* is after the last row of *T*, then let  $T_t$  contain all rows of *T*, preserving their order in *T*.
  - d) If RELATIVE is specified with a zero value of *J*, then
 

Case:

    - i) If the position of *CR* is on a row of *T*, then let  $T_t$  be a sequence of rows comprising that one row.
    - ii) Otherwise, let  $T_t$  be an empty sequence of rows.
- 8) Let *N* be the number of rows in  $T_t$ . If *J* is positive, then let *K* be *J*. If *J* is negative, then let *K* be *N*+*J*+1. If *J* is zero and ABSOLUTE is specified, then let *K* be zero; if *J* is zero and RELATIVE is specified, then let *K* be 1 (one).
- 9) Case:
- a) If *K* is greater than 0 (zero) and not greater than *N*, then *CR* is positioned on the *K*-th row of  $T_t$  and the corresponding row of *T*. That row becomes the current row of *CR*.
  - b) Otherwise, a completion condition is raised: *no data*.
- Case:
- i) If *FO* specifies RELATIVE with *J* equal to 0 (zero), then the position of *CR* is unchanged.
  - ii) If *FO* implicitly or explicitly specifies NEXT, specifies ABSOLUTE or RELATIVE with *K* greater than *N*, or specifies LAST, then *CR* is positioned after the last row.
  - iii) Otherwise, *FO* specifies PRIOR, FIRST, or ABSOLUTE, or RELATIVE with *K* not greater than *N* and *CR* is positioned before the first row.

## Conformance Rules

*None.*

## 15.4 Effect of closing a cursor

### Subclause Signature

```
"Effect of closing a cursor" [General Rules] (
Parameter: "CURSOR",
Parameter: "DISPOSITION"
)
```

### Function

Specify the effect of closing a cursor.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *CR* be the *CURSOR* and let *DISP* be the *DISPOSITION* in an application of this Subclause.
- 2) If *CR* is not in the open state, then an exception condition is raised: *invalid cursor state*.
- 3) Let *RS* be the result set descriptor of *CR*.
- 4) *CR* is placed in the closed state.
- 5) Case:
  - a) If the operational returnability property of *RS* is WITHOUT RETURN, then *RS* is destroyed.
  - b) If *DISP* is DESTROY, then *RS* is removed from the result set sequence that includes *RS*, if any, and *RS* is destroyed.  
NOTE 409 — Otherwise, *RS* is not destroyed.
- 6) If *CR* is a received cursor, and the result set sequence *RSS* of the SQL-invoked routine specified in the cursor declaration descriptor of *CR* is not empty, then the General Rules of Subclause 15.2, "Effect of receiving a result set", are applied with *CR* as *CURSOR* and *RSS* as *RESULT SET SEQUENCE*.

### Conformance Rules

*None.*

## 15.5 Effect of a positioned delete

### Subclause Signature

```
"Effect of a positioned delete" [General Rules] (
Parameter: "CURSOR" ,
Parameter: "STATEMENT" ,
Parameter: "TARGET"
)
```

### Function

Specify the effect of a positioned delete.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *CR* be the *CURSOR*, let *DSP* be the *STATEMENT*, and let *TT* be the *TARGET* in an application of this Subclause.
- 2) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only, and not every leaf generally underlying table of *CR* is a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 3) If there is any sensitive cursor *SCR*, other than *CR*, that is currently open in the SQL-transaction in which *DSP* is being executed, then  
Case:
  - a) If *SCR* has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of *DSP* is made visible to *SCR* or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - b) Otherwise, whether the change resulting from the successful execution of *DSP* is made visible to *SCR* is implementation-defined.
- 4) If there is any insensitive cursor *ICR*, other than *CR*, that is currently open, then either the change resulting from the successful execution of *DSP* is invisible to *ICR*, or an exception condition is raised: *cursor sensitivity exception — request failed*.
- 5) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.

- 6) If *CR* is not positioned on a row, then an exception condition is raised: *invalid cursor state*.
- 7) If *CR* is a holdable cursor and a <fetch statement> has not been issued against *CR* within the current SQL-transaction, then an exception condition is raised: *invalid cursor state*.
- 8) Let *T* be the simply underlying table of *CR* and let *LUT* be the leaf underlying table of *T*.
- 9) Let *R* be the current row of *CR*. Exactly one row *RI* in *LUT* such that each field in *R* is identical to the corresponding field in *RI* is identified for deletion from *LUT*.

NOTE 410 — In case more than one row *RI* satisfies the stated condition, it is implementation-dependent which one is identified for deletion.

NOTE 411 — Identifying a row for deletion is an implementation-dependent mechanism.

- 10) Whether the current row is removed from the sequence of rows of the result set descriptor of *CR* is implementation-defined.
- 11) Case:
  - a) If *LUT* is a base table, then:
    - i) Case:
      - 1) If *TT* specifies ONLY, then *LUT* is *identified for deletion processing without subtables*.
      - 2) Otherwise, *LUT* is *identified for deletion processing with subtables*.

NOTE 412 — Identifying a base table for deletion processing, with or without subtables, is an implementation-dependent mechanism.
    - ii) The General Rules of Subclause 15.7, “Effect of deleting rows from base tables”, are applied.
  - b) If *LUT* is a viewed table, then the General Rules of Subclause 15.9, “Effect of deleting some rows from a viewed table”, are applied with *TT* as *VIEW NAME*.
- 12) If, while *CR* is open, the row from which the current row of *CR* is derived has been marked for deletion by any <delete statement: searched>, <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned>, that identifies any cursor other than *CR*, or updated by any <update statement: searched>, <merge statement>, <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies any cursor other than *CR*, then a completion condition is raised: *warning — cursor operation conflict*.
- 13) If the execution of *DSP* deleted the last row of *CR*, then the position of *CR* is after the last row; otherwise, the position of *CR* is before the next row.

## Conformance Rules

*None.*

## 15.6 Effect of a positioned update

### Subclause Signature

```
"Effect of a positioned update" [General Rules] (
Parameter: "CURSOR" ,
Parameter: "SET CLAUSE LIST" ,
Parameter: "STATEMENT" ,
Parameter: "TARGET"
)
```

### Function

Specify the effect of a positioned update.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *CR* be the *CURSOR*, let *SCL* be the *SET CLAUSE LIST*, let *USP* be the *STATEMENT*, and let *TT* be the *TARGET* in an application of this Subclause.
- 2) If the access mode of the current SQL-transaction or the access mode of the branch of the current SQL-transaction at the current SQL-connection is read-only and not every leaf generally underlying table of *CR* is a temporary table, then an exception condition is raised: *invalid transaction state — read-only SQL-transaction*.
- 3) If there is any sensitive cursor *SCR*, other than *CR*, that is currently open in the SQL-transaction in which *USP* is being executed, then  
Case:
  - a) If *SCR* has not been held into a subsequent SQL-transaction, then either the change resulting from the successful execution of *USP* is made visible to *SCR* or an exception condition is raised: *cursor sensitivity exception — request failed*.
  - b) Otherwise, whether the change resulting from the successful execution of *USP* is made visible to *SCR* is implementation-defined.
- 4) If there is any insensitive cursor *ICR*, other than *CR*, that is currently open, then either the change resulting from the successful execution of *USP* is invisible to *ICR*, or an exception condition is raised: *cursor sensitivity exception — request failed*.



- 5) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
  - 6) If *CR* is not positioned on a row, then an exception condition is raised: *invalid cursor state*.
  - 7) If *CR* is a holdable cursor and a <fetch statement> has not been issued against *CR* within the current SQL-transaction, then an exception condition is raised: *invalid cursor state*.
  - 8) An object row is any row of a base table from which the current row of *CR* is derived.
  - 9) If, while *CR* is open, an object row has been marked for deletion by any <delete statement: searched>, <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned> that identifies any cursor other than *CR*, or updated by any <update statement: searched>, <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies any cursor other than *CR*, or or updated by any <merge statement>, then a completion condition is raised: *warning — cursor operation conflict*.
  - 10) The value associated with DEFAULT is the default value for the <object column> in the containing <set clause> contained in *SCL*, as indicated in the General Rules of Subclause 11.5, “<default clause>”.
  - 11) Each <update source> contained in *SCL* is effectively evaluated for the current row before any of the current row's object rows is updated.
  - 12) *CR* remains positioned on its current row, even if an exception condition is raised during evaluation of any <update source>.
  - 13) A candidate new row is constructed by copying the current row of *CR* and updating it as specified by each <set clause> contained in *SCL* by applying the General Rules of Subclause 14.15, “<set clause list>”.
- NOTE 413 — The data values allowable in an object row may be constrained by a WITH CHECK OPTION constraint. The effect of a WITH CHECK OPTION constraint is defined in the General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”.
- 14) Let *T* be the simply underlying table of *CR* and let *LUT* be the leaf underlying table of *T*.
  - 15) Let *CL* be the columns of *T* identified by the <object column>s contained in *SCL*.
  - 16) Let *R1* be the candidate new row and let *R* be the current row of *CR*. Exactly one row *TR* in *T*, such that the value of each field in *R* that is derived from one or more fields in *TR* is identical to the corresponding value that is derived from the same one or more fields in *TR* is identified for replacement in *T*. The current row *R* of *CR* is replaced by *R1*. Let *TR1* be a row consisting of the fields of *R1* and the fields of *TR* that have no corresponding fields in *R1*, ordered according to the order of their corresponding columns in *T*. *TR1* is the replacement row for *TR* and { ( *TR*, *TR1* ) } is the replacement set for *T*.

NOTE 414 — In case more than one row *R1* satisfies the stated condition, it is implementation-dependent which one is identified for replacement.

NOTE 415 — Identifying a row for replacement, associating a replacement row with an identified row, and associating a replacement set with a table are implementation-dependent mechanisms.

- 17) Case:
  - a) If *LUT* is a base table, then:
    - i) Case:
      - 1) If *TT* specifies ONLY, then *LUT* is identified for replacement processing without subtables with respect to object columns *CL*.

## 15.6 Effect of a positioned update

- 2) Otherwise, *LUT* is identified for replacement processing with subtables with respect to object columns *CL*.

NOTE 416 — Identifying a base table for replacement processing, with or without subtables, is an implementation-dependent mechanism. In general, though not here, the list of object columns can be empty.

- ii) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.
- b) If *LUT* is a viewed table, then General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”, are applied with *TT* as *VIEW NAME* and the replacement set for *T* as *REPLACEMENT SET FOR VIEW NAME*.

## Conformance Rules

*None.*

## 15.7 Effect of deleting rows from base tables

*This Subclause is modified by Subclause 15.1, “Effect of deleting rows from base tables”, in ISO/IEC 9075-9.*

### Function

Specify the effect of deleting rows from one or more base tables.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let  $TT$  be the set consisting of every base table that is identified for deletion processing. Let  $S$  be the set consisting of every row identified for deletion in some table in  $TT$ .
- 2) For every row  $R$  in  $S$ , every row  $SR$  that is a subrow or a superrow of  $R$  is identified for deletion from the base table  $BT$  containing  $SR$ , and  $BT$  is identified for deletion processing.
- 3) Let  $SSC$  be the set of state changes in the current trigger execution context.
- 4) For every table  $T$  in  $TT$ , for every table  $ST$  that is a supertable of  $T$  or, unless  $T$  is identified for deletion processing without subtables, a subtable of  $T$ ,

Case:

- a) If a state change  $SC$  exists in  $SSC$  with subject table  $ST$  and trigger event DELETE, then one copy each of every row of  $ST$  that is identified for deletion in  $ST$  is added to the set of transitions of  $SC$ .
- b) Otherwise, a state change  $SC$  is added to  $SSC$  as follows:
  - i) The set of transitions of  $SC$  consists of one copy each of every row of  $ST$  that is identified for deletion in  $ST$ .
  - ii) The trigger event of  $SC$  is DELETE.
  - iii) The subject table of  $SC$  is  $ST$ .
  - iv) The column list of  $SC$  is empty.
  - v) The set of statement-level triggers for which  $SC$  is considered as executed is empty.
  - vi) The set of row-level triggers consists of each row-level trigger that is activated by  $SC$ , paired with the empty set (of rows considered as executed).
- 5) The Syntax Rules and General Rules of Subclause 15.16, “Execution of BEFORE triggers”, are applied with  $SSC$  as the *SET OF STATE CHANGES*.

- 6) 09 Every row that is identified for deletion in some table identified for deletion processing is marked for deletion. These rows are no longer identified for deletion, nor are their containing tables identified for deletion processing.

NOTE 417 — “Marking for deletion” is an implementation-dependent mechanism.

## **Conformance Rules**

*None.*

## 15.8 Effect of deleting some rows from a derived table

### Subclause Signature

"Effect of deleting some rows from a derived table" [General Rules] (  
 Parameter: "TABLE"  
 )

### Function

Specify the effect of deleting some rows from a derived table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *QE* be the *TABLE* in an application of this Subclause.
- 2) Let *T* be the result of evaluating *QE*.
- 3) Case:
  - a) If *QE* simply contains a <query primary> that immediately contains a <query expression body>, then let *QEB* be that <query expression body>. The General Rules of this Subclause are applied with the table identified by *QEB* as *TABLE*.
  - b) If *QE* simply contains a <query expression body> *QEB* that specifies UNION ALL, then let *LO* and *RO* be the <query expression body> and the <query term>, respectively, that are immediately contained in *QEB*. Let *T1* and *T2* be the tables identified by *LO* and *RO*, respectively.
    - i) For every row *R* in *T* that has been identified for deletion, let *RD* be the row in either *T1* or *T2* from which *R* has been derived and let *TD* be that table. Identify *RD* for deletion.
    - ii) The General Rules of this Subclause are applied with *LO* as *TABLE*.
    - iii) The General Rules of this Subclause are applied with *RO* as *TABLE*.
  - c) Otherwise, let *QS* be the <query specification> simply contained in *QE*. Let *TE* be the <table expression> immediately contained in *QS*, and *TREF* be the <table reference>s simply contained in the <from clause> of *TE*.
    - i) Case:

## 15.8 Effect of deleting some rows from a derived table

- 1) If *TREF* contains only one <table reference>, then let  $TR_1$  be that <table reference>, and let  $m$  be 1 (one).
- 2) Otherwise, let  $m$  be the number of <table reference>s that identify tables with respect to which *QS* is one-to-one. Let  $TR_i$ ,  $1 (one) \leq i \leq m$ , be those <table reference>s.

NOTE 418 — The notion of one-to-one <query specification>s is defined in Subclause 7.12, “<query specification>”.

- ii) Let  $TT_i$ ,  $1 (one) \leq i \leq m$ , be the table identified by  $TR_i$ .
- iii) For every row  $R$  of  $T$  that has been identified for deletion, and for  $i$  ranging from 1 (one) to  $m$ , let  $RD$  be the row in  $TT_i$  from which  $R$  has been derived. Identify that  $RD$  for deletion.
- iv) For  $i$  ranging from 1 (one) to  $m$ ,  
Case:
  - 1) If  $TT_i$  is a base table, then  
Case:
    - A) If  $TR_i$  specifies ONLY, then  $TT_i$  is identified for deletion processing without subtables.
    - B) Otherwise,  $TT_i$  is identified for deletion processing with subtables.
  - 2) If  $TT_i$  is a viewed table, then the General Rules of Subclause 15.9, “Effect of deleting some rows from a viewed table”, are applied with the <table name> of  $TT_i$  as *VIEW NAME*.
  - 3) Otherwise, the General Rules of this Subclause are applied with  $TR_i$  as *TABLE*.
- v) The General Rules of Subclause 15.7, “Effect of deleting rows from base tables”, are applied.

## Conformance Rules

*None.*

## 15.9 Effect of deleting some rows from a viewed table

### Subclause Signature

"Effect of deleting some rows from a viewed table" [General Rules] (  
 Parameter: "VIEW NAME"  
 )

### Function

Specify the effect of deleting some rows from a viewed table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *VN* be the *VIEW NAME* in an application of this Subclause.
- 2) Let *V* be the viewed table identified by *VN*.
- 3) Case:
  - a) If *V* is trigger deletable, then:
    - i) Let *TR* be the delete *INSTEAD OF* trigger whose subject table is *V*.
    - ii) A state change *SC* is created as follows:
      - 1) The set of transitions of *SC* consists of one copy of every row of *V* that is identified for deletion.
      - 2) The trigger event of *SC* is *DELETE*.
      - 3) The subject table of *SC* is *V*.
      - 4) The column list of *SC* is empty.
      - 5) The set of statement-level triggers for which *SC* is considered as executed is empty.
      - 6) The set of row-level triggers consists of *TR* paired with the empty set (of rows considered as executed).
    - iii) The General Rules of Subclause 15.19, "Execution of triggers", are applied with *TR* as *TRIGGER* and *SC* as *STATE CHANGE*.
  - b) Otherwise,

**15.9 Effect of deleting some rows from a viewed table**

- i) If *VN* specifies ONLY, then let *QE* be the original <query expression> included in the descriptor of the view *V* identified by *VN*; otherwise, let *QE* be the <query expression> contained in that descriptor. Let *T* be the result of evaluating *QE*.
- ii) For each row *R* of *V* that has been identified for deletion, let *RD* be the row in *T* from which *R* has been derived; identify that row for deletion.
- iii) The General Rules of [Subclause 15.8](#), “Effect of deleting some rows from a derived table”, are applied with *QE* as *TABLE*.

**Conformance Rules**

*None.*



## 15.10 Effect of inserting tables into base tables

This Subclause is modified by Subclause 15.2, “Effect of inserting tables into base tables”, in ISO/IEC 9075-9.

### Function

Specify the effect of inserting each of one or more given tables into its associated base table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *SSC* be the set of state changes in the current trigger execution context.
- 2) For each base table *T* that is identified for insertion, let *S* be the source table for *T*.
  - a) If some column *IC* of *T* is the identity column of *T*, then for each row in *S* whose site *ICS* corresponding to *IC* is marked as *unassigned*:
    - i) *ICS* is no longer marked as unassigned.
    - ii) Let *NV* be the result of applying the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”, with the sequence generator descriptor included in the column descriptor of *IC* as *SEQUENCE*.  
Case:
      - 1) If the declared type of *IC* is a distinct type *DIST*, then let *ICNV* be *DIST(NV)*.
      - 2) Otherwise, let *ICNV* be *NV*.
    - iii) The General Rules of Subclause 9.2, “Store assignment”, are applied with *ICS* as *TARGET* and *ICNV* as *VALUE*.
  - b) Every proper supertable *ST* of *T* is identified for insertion. A source table for insertion into each *ST* is constructed as follows:
    - i) Let *S* be the source table for the insertion into *T*. Let *TVC* be some <table value constructor> whose value is *S*.
    - ii) Let *n* be the number of column descriptors included in the table descriptor of *ST* and let *CD<sub>i</sub>*, 1 (one) ≤ *i* ≤ *n*, be those column descriptors. Let *SL* be a <select list> containing *n* <select sublist>s such that, for *i* ranging from 1 (one) to *n*, the *i*-th <select sublist> consists of the column name included in *CD<sub>i</sub>*.

## 15.10 Effect of inserting tables into base tables

- iii) The source table for insertion into *ST* consists of the rows in the result of the <query expression>:

```
SELECT SL FROM TVC
```

- 3) For every base table *BT* that is identified for insertion,

Case:

- a) If a state change *SC* exists in *SSC* with subject table *ST* and trigger event INSERT, then the rows in the source table for *BT* are added to the set of transitions of *SC*.
- b) Otherwise, a state change *SC* is added to *SSC* as follows:
  - i) The set of transitions of *SC* consists of the rows in the source table for *BT*.
  - ii) The trigger event of *SC* is INSERT.
  - iii) The subject table of *SC* is *BT*.
  - iv) The column list of *SC* is empty.
  - v) The set of statement-level triggers for which *SC* is considered as executed is empty.
  - vi) The set of row-level triggers consists of each row-level trigger that is activated by *SC*, paired with the empty set (of rows considered as executed).
- 4) The Syntax Rules and General Rules of [Subclause 15.16, “Execution of BEFORE triggers”](#), are applied with *SSC* as the *SET OF STATE CHANGES*.
- 5) For every state change *SC* in *SSC*, let *SOT* be the set of transitions in *SC* and let *BT* be the subject table of *SC*.
  - a) In each row *R* in *SOT*, for each site *GCS* in *R* corresponding to a generated column *GC*, let *GCR* be the result of evaluating, for *R*, the generation expression included in the column descriptor of *GC*. The General Rules of [Subclause 9.2, “Store assignment”](#), are applied with *GCS* as *TARGET* and *GCR* as *VALUE*.
  - b) 09 Every row in *SOT* is inserted into *BT* and *BT* is no longer identified for insertion.

## Conformance Rules

*None.*

## 15.11 Effect of inserting a table into a derived table

### Subclause Signature

"Effect of inserting a table into a derived table" [General Rules] (  
 Parameter: "SOURCE",  
 Parameter: "TARGET"  
 )

### Function

Specify the effect of inserting a table into a derived table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let  $Q$  be the *SOURCE* and let  $T$  be the *TARGET* in an application of this Subclause.
- 2) Let  $QE$  be the <query expression> included in the descriptor of  $T$ .

Case:

- a) If  $QE$  simply contains a <query primary> that immediately contains a <query expression body>, then let  $QEB$  be that <query expression body>. The General Rules of this Subclause are applied with  $Q$  as *SOURCE* and the result of  $QEB$  as *TARGET*.
- b) Otherwise, let  $QS$  be the <query specification> simply contained in  $QE$ . Let  $TE$  be the <table expression> immediately contained in  $QS$ , and  $TREF$  be the <table reference>s simply contained in the <from clause> of  $TE$ . Let  $SL$  be the <select list> immediately contained in  $QS$ , and  $n$  the number of <value expression>s  $VE_j$ ,  $1 \text{ (one)} \leq j \leq n$ , simply contained in  $SL$ .
  - i) Case:
    - 1) If  $TREF$  contains only one <table reference>, then let  $TR_1$  be that <table reference>, and let  $m$  be 1 (one).
    - 2) Otherwise, let  $m$  be the number of <table reference>s that identify tables with respect to which  $QS$  is one-to-one. Let  $TR_i$ ,  $1 \text{ (one)} \leq i \leq m$ , be those <table reference>s.
  - ii) Let  $TT_i$ ,  $1 \text{ (one)} \leq i \leq m$ , be the table identified by  $TR_i$ , and let  $S_i$  be an initially empty table of candidate rows for  $TT_i$ .

**15.11 Effect of inserting a table into a derived table**

- iii) For every row  $R$  of  $Q$ , and for  $i$  ranging from 1 (one) to  $m$ :
  - 1) A candidate row of  $TT_i$  is effectively created in which the value of each column is its default value, as specified the General Rules of Subclause 11.5, “<default clause>”. The candidate row includes every column of  $TT_i$ .
  - 2) For  $j$  ranging from 1 (one) to  $n$ , let  $C$  be a column of some candidate row identified by  $VE_j$ , and let  $SV$  be the  $j$ -th value of  $R$ . The General Rules of Subclause 9.2, “Store assignment”, are applied to  $C$  and  $SV$  as *TARGET* and *SOURCE*, respectively.
  - 3) The candidate row is added to the corresponding  $S_i$ .
- iv) For  $i$  ranging from 1 (one) to  $m$ ,
 

Case:

  - 1) If  $TT_i$  is a base table, then  $TT_i$  is identified for insertion of source table  $S_i$ .
  - 2) If  $TT_i$  is a viewed table, the General Rules of Subclause 15.12, “Effect of inserting a table into a viewed table”, are applied with  $S_i$  as *SOURCE* and  $TT_i$  as *TARGET*.
  - 3) Otherwise, the General Rules of Subclause 15.11, “Effect of inserting a table into a derived table”, are applied with  $S_i$  as *SOURCE* and  $TT_i$  as *TARGET*.
- v) The General Rules of Subclause 15.10, “Effect of inserting tables into base tables”, are applied.

**Conformance Rules**

*None.*

## 15.12 Effect of inserting a table into a viewed table

### Subclause Signature

```
"Effect of inserting a table into a viewed table" [General Rules] (
Parameter: "SOURCE",
Parameter: "TARGET"
)
```

### Function

Specify the effect of inserting a table into a viewed table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *S* be the *SOURCE* and let *T* be the *TARGET* in an application of this Subclause.
- 2) Let *TD* be the view descriptor of *T*. Let *QE* be the original <query expression> included in *TD*.
- 3) Case:
  - a) If *T* is trigger insertable-into, then:
    - i) Let *TR* be the insert INSTEAD OF trigger whose subject table is *T*.
    - ii) A state change *SC* is created as follows:
      - 1) The set of transitions of *SC* consists of one copy of every row of *S*.
      - 2) The trigger event of *SC* is INSERT.
      - 3) The subject table of *SC* is *T*.
      - 4) The column list of *SC* is empty.
      - 5) The set of statement-level triggers for which *SC* is considered as executed is empty.
      - 6) The set of row-level triggers consists of *TR* paired with the empty set (of rows considered as executed).
    - iii) The General Rules of Subclause 15.19, "Execution of triggers", are applied with *TR* as *TRIGGER* and *SC* as *STATE CHANGE*.
  - b) Otherwise,

## 15.12 Effect of inserting a table into a viewed table

Case:

i) If *TD* indicates WITH CHECK OPTION, then:

1) Case:

A) If *TD* specifies LOCAL, then let *VD* be a view descriptor derived from *TD* by removing the WITH CHECK OPTION indication.

B) Otherwise, let *VD* be a view descriptor derived from *TD* as follows:

I) The WITH CHECK OPTION indication is removed.

II) Every reference contained in *QE* to an underlying table *UV* of *QE* that is a viewed table is replaced by a reference to a view whose descriptor is identical to that of *UV* except that WITH CASCADED CHECK OPTION is indicated.

2) The General Rules of this Subclause are applied with *S* as *SOURCE* and the view described by *VD* as *TARGET*.

3) If the result of

```
EXISTS (SELECT * FROM S
 EXCEPT ALL
 SELECT * FROM T)
```

is *True*, then an exception condition is raised: *with check option violation*.

ii) Otherwise, the General Rules of Subclause 15.11, “Effect of inserting a table into a derived table”, are applied with *S* as *SOURCE* and *T* as *TARGET*.

## Conformance Rules

*None.*

## 15.13 Effect of replacing rows in base tables

*This Subclause is modified by Subclause 15.3, “Effect of replacing rows in base tables”, in ISO/IEC 9075-9.*

### Function

Specify the effect of replacing some of the rows in one or more base tables.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let  $TT$  be the set consisting of every base table that is identified for replacement processing. Let  $S$  be the set consisting of every row identified for replacement in every table in  $TT$ .
- 2) For every base table  $T$  in  $TT$ , let  $OC$  be the set consisting of every object column with respect to which  $T$  is identified for replacement processing and every generated column of  $T$  that depends on at least one of these object columns. Every table  $ST$  that is a subtable or supertable of  $T$  is identified for replacement processing with respect to the intersection (possibly empty) of  $OC$  and the columns of  $ST$ .
- 3) For every row  $R$  that is identified for replacement in some table  $T$  in  $TT$ , every row  $SR$  that is a subrow or a superrow of  $R$  is identified for replacement in the base table  $ST$  that contains  $SR$ . The replacement set  $RST$  for  $ST$  is derived from the replacement set  $RR$  for  $T$  as follows.

Case:

- a) If  $ST$  is a subtable of  $T$ , then each replacement row in  $RST$  is the corresponding replacement row in  $RR$  extended with those fields of the corresponding identified row in  $ST$  that have no corresponding column in  $T$ .
  - b) If  $ST$  is a supertable of  $T$ , then each replacement row in  $RST$  is the corresponding replacement row in  $RR$  minus those fields that have no corresponding column in  $ST$ .
- 4) Let  $SSC$  be the set of state changes in the current trigger execution context.
  - 5) For every table  $ST$  that is identified for replacement processing, let  $TL$  be the set consisting of the names of the columns of  $ST$ . For every subset  $STL$  of  $TL$  such that either  $STL$  is empty or the intersection of  $STL$  and  $OC$  is not empty:
    - a) If some column  $IC$  of  $T$  is the identity column of  $ST$ , then, for each row identified for replacement in  $ST$  whose site  $ICS$  corresponding to  $IC$  is marked as *unassigned*:
      - i) Let  $NV$  be the result of applying the General Rules of Subclause 9.23, “Generation of the next value of a sequence generator”, with the sequence generator descriptor included in the column descriptor of  $IC$  as *SEQUENCE*.

## 15.13 Effect of replacing rows in base tables

Case:

- 1) If the declared type of *IC* is a distinct type *DIST*, then let *ICNV* be *DIST(NV)*.
  - 2) Otherwise, let *ICNV* be *NV*.
- ii) The General Rules of [Subclause 9.2, “Store assignment”](#), are applied with *ICS* as *TARGET* and *ICNV* as *VALUE*.
- b) All sites in *ST* that are marked as unassigned cease to be so marked.
- c) Case:
- i) If a state change *SC* exists in *SSC* with subject table *ST*, trigger event *UPDATE*, and column list *STL*, then the row pairs formed by pairing each row identified for replacement in *ST* with its corresponding replacement row are added to the set of transitions of *SC*.
  - ii) Otherwise, a state change *SC* is added to *SSC* as follows:
    - 1) The set of transitions of *SC* consists of row pairs formed by pairing each row identified for replacement in *ST* with its corresponding replacement row.
    - 2) The trigger event of *SC* is *UPDATE*.
    - 3) The subject table of *SC* is *ST*.
    - 4) The column list of *SC* is *STL*.
    - 5) The set of statement-level triggers for which *SC* is considered as executed is empty.
    - 6) The set of row-level triggers consists of each row-level trigger that is activated by *SC*, paired with the empty set (of rows considered as executed).
- 6) The Syntax Rules and General Rules of [Subclause 15.16, “Execution of BEFORE triggers”](#), are applied with *SSC* as the *SET OF STATE CHANGES*.
- 7) For each set of transitions *RST* in each state change *SC* in *SSC*, in each row *R* in *RST*, for each site *GCS* in *R* corresponding to a generated column *GC* in the subject table of *SC*, let *GCR* be the result of evaluating, for *R*, the generation expression included in the column descriptor of *GC*. The General Rules of [Subclause 9.2, “Store assignment”](#), are applied with *GCS* as *TARGET* and *GCR* as *VALUE*.
- 8) 09 For every table *T* in *TT*, for every table *ST* that is a supertable or a subtable of *T*, for every row *R* that is identified for replacement in *ST*, *R* is replaced by its new transition variable. *R* is no longer identified for replacement. *ST* is no longer identified for replacement processing.

## Conformance Rules

*None.*



## 15.14 Effect of replacing some rows in a derived table

### Subclause Signature

```
"Effect of replacing some rows in a derived table" [General Rules] (
Parameter: "TABLE" ,
Parameter: "REPLACEMENT SET FOR TABLE"
)
```

### Function

Specify the effect of replacing some rows in a derived table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *QE* be the *TABLE* and let *RS* be the *REPLACEMENT SET FOR TABLE* in an application of this Subclause.
- 2) Let *T* be the result of evaluating *QE*. Let *CL* be the object columns of *QE*.
- 3) Case:
  - a) If *QE* simply contains a <query primary> that immediately contains a <query expression body>, then let *QEB* be that <query expression body>. The General Rules of this Subclause are applied with *QEB* as *TABLE* and *RS* as *REPLACEMENT SET FOR TABLE*.
  - b) If *QE* simply contains a <query expression body> *QEB* that specifies UNION ALL, let *LO* and *RO* be the <query expression body> and the <query term>, respectively, that are immediately contained in *QEB*. Let *T1* and *T2* be the tables identified by *LO* and *RO*, respectively. Let the columns of *T1* and *T2* that are underlying columns of the object columns of *CL* be the object columns *CL1* and *CL2*, respectively. Let *RS1* and *RS2* be the initially empty replacement sets for *T1* and *T2*, respectively.
    - i) For every pair (*SR*, *CNR*) of *RS*,
 

Case:

      - 1) If *SR* has been derived from a row of *T1*, then identify that row *SR1* for replacement by *CNR*; the pair (*SR1*, *CNR*) is effectively added to *RS1*.
      - 2) Otherwise, let *SR2* be the row of *T2* from which *SR* has been derived; identify that row for replacement by *CNR*; the pair (*SR2*, *CNR*) is effectively added to *RS2*.

## 15.14 Effect of replacing some rows in a derived table

- ii) The General Rules of this Subclause are applied with *LO* as *TABLE* and *RS1* as *REPLACEMENT SET FOR TABLE*.
  - iii) The General Rules of this Subclause are applied with *RO* as *TABLE* and *RS2* as *REPLACEMENT SET FOR TABLE*.
- c) Otherwise, let *QS* be the <query specification> simply contained in *QE*. Let *TE* be the <table expression> immediately contained in *QS*, and let *TREF* be the <table reference>s simply contained in the <from clause> of *TE*. Let *SL* be the <select list> immediately contained in *QS*, and let *n* be the number of <value expression>s *VE<sub>j</sub>*,  $1 \text{ (one)} \leq j \leq n$ , simply contained in *SL*.
- i) Case:
    - 1) If *TREF* contains only one <table reference>, then let *TR<sub>1</sub>* be that <table reference>, and let *m* be 1 (one).
    - 2) Otherwise, let *m* be the number of <table reference>s that identify tables with respect to which *QS* is one-to-one. Let *TR<sub>i</sub>*,  $1 \text{ (one)} \leq i \leq m$ , be those <table reference>s.
  - ii) Let *TT<sub>i</sub>*,  $1 \text{ (one)} \leq i \leq m$ , be the table identified by *TR<sub>i</sub>*, let *RS<sub>i</sub>* be an initially empty replacement set for *TT<sub>i</sub>*, and let *CL<sub>i</sub>* be the object column list of *TT<sub>i</sub>*, such that every column of *CL<sub>i</sub>* is an underlying column of *CL*.
  - iii) For every pair (*SR*, *CNR*) of *RS*, and for *i* ranging from 1 (one) to *m*:
    - 1) Let *SRTI* be the row of *TT<sub>i</sub>* from which *SR* has been derived.
    - 2) A candidate row *CNRI* of *TT<sub>i</sub>* is effectively created in which the value of each column is its default value, as specified the General Rules of Subclause 11.5, “<default clause>”. The candidate row includes every column of *TT<sub>i</sub>*.
    - 3) For *j* ranging from 1 (one) to *n*, let *C* be a column of some candidate row identified by *VE<sub>j</sub>*, and let *SV* be the *j*-th value of *CNR*. The General Rules of Subclause 9.2, “Store assignment”, are applied to *C* and *SV* as *TARGET* and *SOURCE*, respectively.
    - 4) Identify *SRTI* for replacement by *CNRI*; the pair (*SRTI*, *CNRI*) is effectively added to *RS<sub>i</sub>*.
  - iv) For *i* ranging from 1 (one) to *m*
    - Case:
      - 1) If *TT<sub>i</sub>* is a base table, then
        - Case:
          - A) If *TR<sub>i</sub>* specifies ONLY, then *TT<sub>i</sub>* is identified for replacement processing without subtables with respect to the object columns *CL<sub>i</sub>*.
          - B) Otherwise, *TT<sub>i</sub>* is identified for replacement processing with subtables with respect to the object columns *CL<sub>i</sub>*.
      - 2) If *TT<sub>i</sub>* is a viewed table, then the General Rules of Subclause 15.15, “Effect of replacing some rows in a viewed table”, are applied with the <table name> of *TT<sub>i</sub>* as *VIEW NAME* and *RS<sub>i</sub>* as *REPLACEMENT SET FOR VIEW NAME*.

**15.14 Effect of replacing some rows in a derived table**

- 3) If  $TT_i$  is a derived table, then the General Rules of this Subclause are applied with  $TR_i$  as *TABLE* and  $RS_i$  as *REPLACEMENT SET FOR TABLE*.
- v) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.

**Conformance Rules**

*None.*

## 15.15 Effect of replacing some rows in a viewed table

### Subclause Signature

```
"Effect of replacing some rows in a viewed table" [General Rules] (
Parameter: "VIEW NAME",
Parameter: "REPLACEMENT SET FOR VIEW NAME"
)
```

### Function

Specify the effect of replacing some rows in a viewed table.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *VN* be the *VIEW NAME* and let *RS* be the *REPLACEMENT SET FOR VIEW NAME* in an application of this Subclause.
- 2) Let *T* be the view identified by *VN*. Let *TD* be the view descriptor of *T*. If *VN* specifies *ONLY*, then let *QE* be the original <query expression> included in *TD*; otherwise, let *QE* be the <query expression> included in *TD*.
- 3) Case:
  - a) If *T* is trigger updatable, then:
    - i) Let *TR* be the update *INSTEAD OF* trigger whose subject table is *T*.
    - ii) A state change *SC* is created as follows:
      - 1) The set of transitions of *SC* consists of copies of row pairs in the replacement set for *T*.
      - 2) The trigger event of *SC* is *UPDATE*.
      - 3) The subject table of *SC* is *T*.
      - 4) The column list of *SC* is empty.
      - 5) The set of statement-level triggers for which *SC* is considered as executed is empty.
      - 6) The set of row-level triggers consists of *TR* paired with the empty set (of rows considered as executed).

## 15.15 Effect of replacing some rows in a viewed table

- iii) The General Rules of Subclause 15.19, “Execution of triggers”, are applied with *TR* as *TRIGGER* and *SC* as *STATE CHANGE*.
- b) Otherwise,
  - Case:
    - i) If *TD* indicates WITH CHECK OPTION, then:
      - 1) Case:
        - A) If *TD* specifies LOCAL, then let *VD* be a view descriptor derived from *TD* by removing the WITH CHECK OPTION indication.
        - B) Otherwise, let *VD* be a view descriptor derived from *TD* as follows:
          - I) The WITH CHECK OPTION indication is removed.
          - II) Every reference contained in *QE* to an underlying table *UV* of *QE* that is a viewed table is replaced by a reference to a view whose descriptor is identical to that of *UV* except that WITH CASCADED CHECK OPTION is indicated.
      - 2) The General Rules of this Subclause are applied with the <table name> *V* of the view described by *VD* as *VIEW NAME* and *RS* as *REPLACEMENT SET FOR VIEW NAME*.
      - 3) Let *S* be the table consisting of the candidate new rows of *RS*. If the result of
 

```

EXISTS (SELECT * FROM S
 EXCEPT ALL
 SELECT * FROM T)

```

 is *True*, then an exception condition is raised: *with check option violation*.
    - ii) Otherwise, the General Rules of Subclause 15.14, “Effect of replacing some rows in a derived table”, are applied with *QE* as *TABLE* and *RS* as *REPLACEMENT SET FOR TABLE*.

## Conformance Rules

*None.*

## 15.16 Execution of BEFORE triggers

### Function

Define the execution of BEFORE triggers.

### Syntax Rules

- 1) Let *SSC* be the *SET OF STATE CHANGES* specified in an application of this Subclause.
- 2) Let *BT* be the set of BEFORE triggers that are activated by some state change in *SSC*.  
NOTE 419 — Activation of triggers is defined in Subclause 4.38, “Triggers”.
- 3) Let *NT* be the number of triggers in *BT* and let  $TR_k$  be the  $k$ -th such trigger, ordered according to their order of execution. Let  $SC_k$  be the state change in *SSC* that activated  $TR_k$ .

NOTE 420 — Ordering of triggers is defined in Subclause 4.38, “Triggers”.

### Access Rules

*None.*

### General Rules

- 1) For  $k$  ranging from 1 (one) to *NT*, apply the General Rules of Subclause 15.19, “Execution of triggers”, with  $TR_k$  as *TRIGGER* and  $SC_k$  as *STATE CHANGE*, respectively.

### Conformance Rules

*None.*

## 15.17 Execution of referential actions

### Subclause Signature

"Execution of referential actions" [General Rules] (  
Parameter: "CONSTRAINT"  
)

### Function

Specify the effect of referential actions.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *RC* be the *CONSTRAINT* in an application of this Subclause.
- 2) *RC* is a referential constraint.
- 3) Let *M* be the <match type> of *RC* (either SIMPLE, PARTIAL, or FULL).
- 4) Let *UR* be the <update rule> of *RC* and let *DR* be the <delete rule> of *RC*.
- 5) Let *F* be the referencing table of *RC*.
- 6) Case:
  - a) If *M* specifies SIMPLE or FULL, then for a given row in the referenced table, every row that is a subrow or a superrow of a row *R* in *F* such that the referencing column values equal the corresponding referenced column values in *R* for the referential constraint is a *matching row*.
  - b) If *M* specifies PARTIAL, then:
    - i) For a given row in the referenced table, every row that is a subrow or a superrow of a row *R* in *F* such that *R* has at least one non-null referencing column value and the non-null referencing column values of *R* equal the corresponding referenced column values for the referential constraint is a *matching row*.
    - ii) For a given row in the referenced table, every matching row for that given row that is a matching row only to the given row in the referenced table for the referential constraint is a *unique matching row*. For a given row in the referenced table, a matching row for that given row that is not a unique matching row for that given row for the referential constraint is a *non-unique matching row*.

- 7) For each row of the referenced table, its matching rows, unique matching rows, and non-unique matching rows are determined immediately prior to the execution of any <SQL procedure statement>. No new matching rows are added during the execution of that <SQL procedure statement>.

The association between a referenced row and a non-unique matching row is dropped during the execution of that SQL-statement if the referenced row is either marked for deletion or updated to a distinct value on any referenced column that corresponds to a non-null referencing column. This occurs immediately after such a mark for deletion or update of the referenced row. Unique matching rows and non-unique matching rows for a referenced row are evaluated immediately after dropping the association between that referenced row and a non-unique matching row.

- 8) Let *CTEC* be the current trigger execution context. Let *SSC* be the set of state changes in *CTEC*. Let *SC<sub>i</sub>* be a state change in *SSC*.
- 9) For every row of the referenced table that is marked for deletion and has not previously been marked for deletion,

Case:

- a) If *M* specifies SIMPLE or FULL, then

Case:

- i) If *DR* specifies CASCADE, then:

- 1) *F* is identified for deletion processing and every matching row in *F* is identified for deletion.
- 2) The General Rules of Subclause 15.7, “Effect of deleting rows from base tables” are applied.

- ii) If *DR* specifies SET NULL, then:

- 1) Each matching row *MR* in *F* is paired with the candidate replacement row *NMR*, formed by copying *MR* and setting each referencing column in the copy to the null value. *MR* is identified for replacement by *NMR* in *F*. The set of (*MR*, *NMR*) pairs is the replacement set for *F*.
- 2) *F* is identified for replacement processing with subtables with respect to the referencing columns.
- 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.

- iii) If *DR* specifies SET DEFAULT, then:

- 1) Each matching row *MR* in *F* is paired with the candidate replacement row *NMR*, formed by copying *MR* and setting each referencing column in the copy to the default value specified in the General Rules of Subclause 11.5, “<default clause>”. *MR* is identified for replacement by *NMR* in *F*. The set of (*MR*, *NMR*) pairs is the replacement set for *F*.
- 2) *F* is identified for replacement processing with subtables with respect to the referencing columns.
- 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.

- iv) If *DR* specifies RESTRICT and there exists some matching row, then an exception condition is raised: *integrity constraint violation — restrict violation*.

NOTE 421 — If *DR* specifies NO ACTION, then no referential delete action is performed.



b) If *M* specifies PARTIAL, then

Case:

i) If *DR* specifies CASCADE, then:

- 1) *F* is identified for deletion processing and every unique matching row in *F* is identified for deletion.
- 2) The General Rules of Subclause 15.7, “Effect of deleting rows from base tables” are applied.

ii) If *DR* specifies SET NULL, then:

- 1) Each unique matching row *UMR* in *F* is paired with the candidate replacement row *NUMR*, formed by copying *UMR* and setting each referencing column in the copy to the null value. *UMR* is identified for replacement by *NUMR* in *F*. The set of (*UMR*, *NUMR*) pairs is the replacement set for *F*.
- 2) *F* is identified for replacement processing with subtables with respect to the referencing columns.
- 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.

iii) If *DR* specifies SET DEFAULT, then:

- 1) Each unique matching row *UMR* in *F* is paired with the candidate replacement row *NUMR*, formed by copying *UMR* and setting each referencing column in the copy to the default value specified in the General Rules of Subclause 11.5, “<default clause>”. *UMR* is identified for replacement by *NUMR* in *F*. The set of (*UMR*, *NUMR*) pairs is the replacement set for *F*.
- 2) *F* is identified for replacement processing with subtables with respect to the referencing columns.
- 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.

iv) If *DR* specifies RESTRICT and there exists some unique matching row, then an exception condition is raised: *integrity constraint violation — restrict violation*.

NOTE 422 — If *DR* specifies NO ACTION, then no referential delete action is performed.

10) If a non-null value of a referenced column *RC* in the referenced table is updated to a value that is distinct from the current value of *RC*, then, for every member *F* of the subtable family of the referencing table

Case:

a) If *M* specifies SIMPLE or FULL, then

Case:

i) If *UR* specifies CASCADE, then:

- 1) Each matching row *MR* in *F* is paired with the candidate replacement row *NMR*, formed by copying *MR* and setting each referencing column in the copy to the new value of that referenced column. *MR* is identified for replacement by *NMR* in *F*. The set of (*MR*, *NMR*) pairs is the replacement set for *F*.
- 2) *F* is identified for replacement processing with subtables with respect to the referencing columns.

- 3) The General Rules of [Subclause 15.13](#), “Effect of replacing rows in base tables”, are applied.
- ii) If *UR* specifies SET NULL, then
- Case:
- 1) If *M* specifies SIMPLE, then:
    - A) Each matching row *MR* in *F* is paired with the candidate replacement row *NMR*, formed by copying *MR* and setting each referencing column in the copy to the null value. *MR* is identified for replacement by *NMR* in *F*. The set of (*MR*, *NMR*) pairs is the replacement set for *F*.
    - B) *F* is identified for replacement processing with subtables with respect to the referencing columns.
    - C) The General Rules of [Subclause 15.13](#), “Effect of replacing rows in base tables”, are applied.
  - 2) If *M* specifies FULL, then:
    - A) Each matching row *MR* in *F* is paired with the candidate replacement row *NMR*, formed by copying *MR* and setting each referencing column in the copy to the null value. *MR* is identified for replacement by *NMR* in *F*. The set of (*MR*, *NMR*) pairs is the replacement set for *F*.
    - B) *F* is identified for replacement processing with subtables with respect to the referencing columns.
    - C) The General Rules of [Subclause 15.13](#), “Effect of replacing rows in base tables”, are applied.
- iii) If *UR* specifies SET DEFAULT, then:
- 1) Each matching row *MR* in *F* is paired with the candidate replacement row *NMR*, formed by copying *MR* and setting each referencing column in the copy to the default value specified in the General Rules of [Subclause 11.5](#), “<default clause>”. *MR* is identified for replacement by *NMR* in *F*. The set of (*MR*, *NMR*) pairs is the replacement set for *F*.
  - 2) *F* is identified for replacement processing with subtables with respect to the referencing columns.
  - 3) The General Rules of [Subclause 15.13](#), “Effect of replacing rows in base tables”, are applied.
- iv) If *UR* specifies RESTRICT and there exists some matching row, then an exception condition is raised: *integrity constraint violation — restrict violation*.

NOTE 423 — If *UR* specifies NO ACTION, then no referential update action is performed.

- b) If *M* specifies PARTIAL, then

Case:

- i) If *UR* specifies CASCADE, then:
  - 1) Each unique matching row *UMR* in *F* that contains a non-null value in the referencing column *C1* in *F* that corresponds to the updated referenced column *C2* is paired with the candidate replacement row *NUMR*, formed by copying *UMR* and setting *C1* in the copy to the new

value  $V$  of  $C2$ , provided that, in all updated rows in the referenced table that formerly had, during the same execution of the same innermost SQL-statement, that unique matching row as a matching row, the values in  $C2$  have all been updated to a value that is not distinct from  $V$ . If this last condition is not satisfied, then an exception condition is raised: *triggered data change violation*.  $UMR$  is identified for replacement by  $NUMR$  in  $F$ . The set of  $(MR, NMR)$  pairs is the replacement set for  $F$ .

NOTE 424 — Because of the Rules of Subclause 8.2, “<comparison predicate>”, on which the definition of “distinct” relies, the values in  $C2$  may have been updated to values that are not distinct, yet are not identical. Which of these non-distinct values is used for the cascade operation is implementation-dependent.

- 2)  $F$  is identified for replacement processing with subtables with respect to the referencing columns.
- 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.
- ii) If  $UR$  specifies SET NULL, then:
  - 1) Each unique matching row  $UMR$  in  $F$  that contains a non-null value in the referencing column in  $F$  is paired with the candidate replacement row  $NUMR$ , formed by copying  $UMR$  and setting that referencing column in the copy to the null value.  $UMR$  is identified for replacement by  $NUMR$  in  $F$ . The set of  $(UMR, NUMR)$  pairs is the replacement set for  $F$ .
  - 2)  $F$  is identified for replacement processing with subtables with respect to the referencing columns.
  - 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.
- iii) If  $UR$  specifies SET DEFAULT, then:
  - 1) Each unique matching row  $UMR$  in  $F$  that contains a non-null value in the referencing column in  $F$  that corresponds with the updated referenced column is paired with the candidate replacement row  $NUMR$ , formed by copying  $UMR$  and setting that referencing column in the copy to the default value specified in the General Rules of Subclause 11.5, “<default clause>”.  $UMR$  is identified for replacement by  $NUMR$  in  $F$ . The set of  $(UMR, NUMR)$  pairs is the replacement set for  $F$ .
  - 2)  $F$  is identified for replacement processing with subtables with respect to the referencing columns.
  - 3) The General Rules of Subclause 15.13, “Effect of replacing rows in base tables”, are applied.
- iv) If  $UR$  specifies RESTRICT and there exists some unique matching row, then an exception condition is raised: *integrity constraint violation — restrict violation*.

NOTE 425 — If  $UR$  specifies NO ACTION, then no referential update action is performed.

- 11) Let  $ISS$  be the innermost SQL-statement being executed.
- 12) If evaluation of these General Rules during the execution of  $ISS$  would cause an update of some site to a value that is distinct from the value to which that site was previously updated during the execution of  $ISS$ , then an exception condition is raised: *triggered data change violation*.
- 13) If evaluation of these General Rules during the execution of  $ISS$  would cause deletion of a row containing a site that is identified for replacement in that row, then an exception condition is raised: *triggered data change violation*.

- 14) If evaluation of these General Rules during the execution of *ISS* would cause either an attempt to update a row that has been deleted by any <delete statement: positioned>, <dynamic delete statement: positioned>, or <preparable dynamic delete statement: positioned> that identifies some open cursor *CR* or that has been updated by any <update statement: positioned>, <dynamic update statement: positioned>, or <preparable dynamic update statement: positioned> that identifies some open cursor *CR*, or an attempt to mark for deletion such a row, then a completion condition is raised: *warning — cursor operation conflict*.
- 15) For every row *RMD* that is marked for deletion, every subrow of *RMD* and every superrow of *RMD* is marked for deletion.
- 16) If any table *T* is the subject table of a state change in *SSC* that has been created or modified during evaluation of the preceding General Rules of this subclause, then, for every referential constraint descriptor *RC2*, General Rules of this Subclause are applied with *RC2* as *CONSTRAINT*.

NOTE 426 — Thus these rules are repeatedly evaluated until no further transitions are generated.

## Conformance Rules

*None.*

## 15.18 Execution of AFTER triggers

### Function

Define the execution of AFTER triggers.

### Syntax Rules

- 1) Let *SSC* be the *SET OF STATE CHANGES* specified in an application of this Subclause.
- 2) Let *AT* be the set of AFTER triggers that are activated by some state change in *SSC*.  
NOTE 427 — Activation of triggers is defined in Subclause 4.38, “Triggers”.
- 3) Let *NT* be the number of triggers in *AT* and let  $TR_k$  be the  $k$ -th such trigger, ordered according to their order of execution. Let  $SC_k$  be the state change in *SSC* that activated  $TR_k$ .

NOTE 428 — Ordering of triggers is defined in Subclause 4.38, “Triggers”.

### Access Rules

*None.*

### General Rules

- 1) For  $k$  ranging from 1 (one) to *NT*, apply the General Rules of Subclause 15.19, “Execution of triggers”, with  $TR_k$  as *TRIGGER* and  $SC_k$  as *STATE CHANGE*, respectively.

### Conformance Rules

*None.*

## 15.19 Execution of triggers

### Subclause Signature

```
"Execution of triggers" [General Rules] (
Parameter: "TRIGGER",
Parameter: "STATE CHANGE"
)
```

### Function

Define the execution of triggers.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *TR* be the *TRIGGER* and let *SC* be the *STATE CHANGE* in an application of this Subclause.
- 2) Let *TA* be the triggered action included in the trigger descriptor of *TR*. Let *TSS* be the <triggered SQL statement> contained in *TA*. Let *TE* be the trigger event of *SC*. Let *ST* be the set of transitions in *SC*.
- 3) *TR* is executed as follows:  
Case:
  - a) If *TR* is a row-level trigger, then, for each transition *T* in *ST* for which *TR* is not considered as executed, *TA* is invoked and *TR* is considered as executed for *T*. The order in which the transitions in *ST* are taken is implementation-dependent.
  - b) If *TR* is not considered as executed for *SC*, then *TA* is invoked once and *TR* is considered as executed for *SC*.
- 4) When *TA* is invoked:
  - a) Case:
    - i) If *TE* is DELETE, then the old transition table for the invocation of *TA* is *ST*. If *TR* is a row-level trigger, then the value of the old transition variable for the execution of *TSS* is *T*.
    - ii) If *TE* is INSERT, then the new transition table for the invocation of *TA* is *ST*. If *TR* is a row-level trigger, then the value of the new transition variable for the invocation of *TA* is *T*.
    - iii) If *TE* is UPDATE, then the old transition table for the invocation of *TA* is the multiset formed by taking the old rows of the transitions in *ST* and the new transition table for the invocation

of *TA* is the multiset formed by taking the new rows of the transitions in *ST*. If *TR* is a row-level trigger, then the value of the old transition variable for the invocation of *TA* is the old row of *T* and the new transition variable for the invocation of *TA* is the new row of *T*.

b) Case:

- i) If *TA* contains a <search condition> *TASC* and the result of evaluating *TASC* is *True*, then *TSS* is executed.
- ii) If *TA* does not contain a <search condition>, then *TSS* is executed.

5) When *TSS* is executed:

- a) The General Rules of Subclause 23.2, “Pushing and popping the diagnostics area stack”, are applied with “PUSH” as *OPERATION* and the diagnostics area stack as *STACK*.
- b) The authorization identifier of the owner of the schema that includes the trigger descriptor of *TR* is pushed onto the authorization stack.
- c) A new savepoint level is established.
- d) Let *N* be the number of <SQL procedure statement>s simply contained in *TSS*. For *i* ranging from 1 (one) to *N*:
  - i) Let *S<sub>i</sub>* be the *i*-th such <SQL procedure statement>.
  - ii) The General Rules of Subclause 13.5, “<SQL procedure statement>”, are evaluated with *S<sub>i</sub>* as the executing statement.
- e) The <SQL procedure statement>s simply contained in *TSS* are effectively executed in the order in which they are specified in *TSS*.
- f) If, before the completion of the execution of any <SQL procedure statement> simply contained in *TSS*, an attempt is made to execute an SQL-schema statement, an SQL-dynamic statement, or an SQL-session statement then an exception condition is raised: *prohibited statement encountered during trigger execution*.
- g) If *TR* is a BEFORE trigger and if, before the completion of the execution of any <SQL procedure statement> simply contained in *TSS*, an attempt is made to execute an SQL-data change statement or an SQL-invoked routine that possibly modifies SQL-data, then an exception condition is raised: *prohibited statement encountered during trigger execution*.
- h) The current savepoint level is destroyed.
 

NOTE 429 — Destroying a savepoint level destroys all existing savepoints that are established at that level.
- i) The General Rules of Subclause 23.2, “Pushing and popping the diagnostics area stack”, are applied with “POP” as *OPERATION* and the diagnostics area stack as *STACK*.
- j) The top cell in the authorization stack is removed.
- k) If the execution of *TSS* is not successful, then an exception condition is raised: *triggered action exception*. The exception condition that caused *TSS* to fail is raised.

NOTE 430 — Raising the exception condition that caused *TSS* to fail enters the exception information into the diagnostics area that was pushed prior to the execution of *TSS*.

## **Conformance Rules**

*None.*



## 16 Control statements

*This Clause is modified by Clause 14, “Control statements”, in ISO/IEC 9075-4.*

*This Clause is modified by Clause 15, “Control statements”, in ISO/IEC 9075-14.*

### 16.1 <call statement>

*This Subclause is modified by Subclause 11.22, “<call statement>”, in ISO/IEC 9075-10.*

#### Function

Invoke an SQL-invoked routine.

#### Format

```
<call statement> ::=
 CALL <routine invocation>
```

#### Syntax Rules

- 1) Let *RI* be the <routine invocation> immediately contained in the <call statement>.
- 2) The Syntax Rules of Subclause 10.4, “<routine invocation>”, are invoked with *RI* as *ROUTINE INVOCATION*, yielding subject routine *SR* and static SQL argument list *SAL*.
- 3) *SR* shall be an SQL-invoked procedure.

#### Access Rules

*None.*

#### General Rules

- 1) The General Rules of Subclause 10.4, “<routine invocation>”, are invoked with *SR* as *ROUTINE INVOCATION* and *SAL* as *STATIC SQL ARG LIST*.

#### Conformance Rules

*None.*

## 16.2 <return statement>

### Function

Return a value from an SQL function.

### Format

```
<return statement> ::=
 RETURN <return value>

<return value> ::=
 <value expression>
 | NULL
```

### Syntax Rules

- 1) <return statement> shall be contained in an SQL routine body that is simply contained in the <routine body> of an <SQL-invoked function> *F*. Let *RDT* be the <returns data type> of the <returns clause> of *F*.
- 2) The <return value> <null specification> is equivalent to the <value expression>:

```
CAST (NULL AS RDT)
```

- 3) Let *VE* be the <value expression> of the <return value> immediately contained in <return statement>.
- 4) The declared type of *VE* shall be assignable to an item of the data type *RDT*, according to the Syntax Rules of [Subclause 9.2, “Store assignment”](#), with *RDT* and *VE* as *TARGET* and *VALUE*, respectively.

### Access Rules

*None.*

### General Rules

- 1) The value of *VE* is the *returned value* of the execution of the SQL routine body of *F*.
- 2) The execution of the SQL routine body of *F* is terminated immediately.

### Conformance Rules

*None.*

## 17 Transaction management

### 17.1 <start transaction statement>

#### Function

Start an SQL-transaction and set its characteristics.

#### Format

```
<start transaction statement> ::=
 START TRANSACTION [<transaction characteristics>]
```

#### Syntax Rules

*None.*

#### Access Rules

*None.*

#### General Rules

- 1) If an SQL-transaction is currently active, then an exception condition is raised: *invalid transaction state — active SQL-transaction*.
- 2) Case:
  - a) If <transaction characteristics> is omitted, then let *TC* be  
  
*ECAM ISOLATION LEVEL ECIL DIAGNOSTICS SIZE ECNC*  
  
where *ECAM*, *ECIL*, and *ECNC* are the access mode, isolation level and number of conditions, respectively, of the enduring transaction characteristics of the current SQL-session.
  - b) Otherwise, let *TC* be the <transaction characteristics>.
- 3) If <number of conditions> is specified and is less than 1 (one), then an exception condition is raised: *invalid condition number*.
- 4) The <set transaction statement>

```
SET TRANSACTION TC
```

is effectively executed.

**17.1 <start transaction statement>**

NOTE 431 — The characteristics of a transaction begun by a <start transaction statement> are as specified here regardless of the characteristics specified by any preceding <set transaction statement>. That is, even if one or more characteristics are omitted by the <start transaction statement>, the defaults specified in the Syntax Rules of this Subclause and of [Subclause 17.3, “<transaction characteristics>”](#), are effective and are not affected by any (preceding) <set transaction statement> in the same SQL-session.

- 5) An SQL-transaction is initiated.

## **Conformance Rules**

- 1) Without Feature T241, “START TRANSACTION statement”, conforming SQL language shall not contain a <start transaction statement>.

## 17.2 <set transaction statement>

This Subclause is modified by *Subclause 11.21*, “<set transaction statement>”, in ISO/IEC 9075-10.

### Function

Set the characteristics of the next SQL-transaction for the SQL-agent.

NOTE 432 — This statement has no effect on any SQL-transactions subsequent to the next SQL-transaction.

### Format

```
<set transaction statement> ::=
SET [LOCAL] TRANSACTION <transaction characteristics>
```

### Syntax Rules

- 1) If LOCAL is specified, then <transaction characteristics> shall not contain <number of conditions>.

### Access Rules

*None.*

### General Rules

- 1) Case:
  - a) If a <set transaction statement> that does not specify LOCAL is executed, then  
Case:
    - i) If an SQL-transaction is currently active, then an exception condition is raised: *invalid transaction state — active SQL-transaction.*
    - ii) If an SQL-transaction is not currently active, then if there are any holdable cursors remaining open from the previous SQL-transaction and the isolation level of the previous SQL-transaction is not the same as the isolation level determined by the <level of isolation>, then an exception condition is raised: *invalid transaction state — held cursor requires same isolation level.*
  - b) If a <set transaction statement> that specifies LOCAL is executed, then:
    - i) If the SQL-implementation does not support SQL-transactions that affect more than one SQL-server, then an exception condition is raised: *feature not supported — multiple server transactions.*
    - ii) If there is no SQL-transaction that is currently active, then an exception condition is raised: *invalid transaction state — no active SQL-transaction for branch transaction.*
    - iii) If there is an active SQL-transaction and there has been a transaction-initiating SQL-statement executed at the current SQL-connection in the context of the active SQL-transaction, then an exception condition is raised: *invalid transaction state — branch transaction already active.*

## 17.2 &lt;set transaction statement&gt;

- iv) If the transaction access mode of the SQL-transaction is read-only and <transaction access mode> specifies READ WRITE, then an exception condition is raised: *invalid transaction state — inappropriate access mode for branch transaction*.
- v) If the isolation level of the SQL-transaction is SERIALIZABLE and <level of isolation> specifies anything except SERIALIZABLE, then an exception condition is raised: *invalid transaction state — inappropriate isolation level for branch transaction*.
- vi) If the isolation level of the SQL-transaction is REPEATABLE READ and <level of isolation> specifies anything except REPEATABLE READ or SERIALIZABLE, then an exception condition is raised: *invalid transaction state — inappropriate isolation level for branch transaction*.
- vii) If the isolation level of the SQL-transaction is READ COMMITTED and <level of isolation> specifies READ UNCOMMITTED, then an exception condition is raised: *invalid transaction state — inappropriate isolation level for branch transaction*.

NOTE 433 — If the isolation level of the SQL-transaction is READ UNCOMMITTED, then any <level of isolation> is permissible.

- 2) If <number of conditions> is specified and is less than 1 (one), then an exception condition is raised: *invalid condition number*.
- 3) If <number of conditions> is specified and is greater than the implementation-dependent maximum value for <number of conditions> *IDMVNC*, then the implicit <number of conditions> is *IDMVNC*, and a completion condition is raised: *warning — invalid number of conditions*.
- 4) Let *TC* be <transaction characteristics>. Let *CSC* be the current SQL-session context.
- 5) If the explicit or implicit <transaction access mode> contains READ ONLY, then the current access mode of *CSC* is set to *read-only*. Otherwise, the current access mode of *CSC* is set to *read-write*.
- 6) The current isolation level of *CSC* is set to an implementation-defined isolation level that will not exhibit any of the phenomena that the explicit or implicit <level of isolation> contained in *TC* would not exhibit, as specified in Table 8, “SQL-transaction isolation levels and the three phenomena”.
- 7) The current condition area limit of *CSC* is set to the explicit or implicit <number of conditions> contained in *TC*.

## Conformance Rules

- 1) Without Feature T251, “SET TRANSACTION statement: LOCAL option”, conforming SQL language shall not contain a <set transaction statement> that immediately contains LOCAL.

## 17.3 <transaction characteristics>

### Function

Specify transaction characteristics.

### Format

```
<transaction characteristics> ::=
 [<transaction mode> [{ <comma> <transaction mode> }...]]

<transaction mode> ::=
 <isolation level>
 | <transaction access mode>
 | <diagnostics size>

<transaction access mode> ::=
 READ ONLY
 | READ WRITE

<isolation level> ::=
 ISOLATION LEVEL <level of isolation>

<level of isolation> ::=
 READ UNCOMMITTED
 | READ COMMITTED
 | REPEATABLE READ
 | SERIALIZABLE

<diagnostics size> ::=
 DIAGNOSTICS SIZE <number of conditions>

<number of conditions> ::=
 <simple value specification>
```

### Syntax Rules

- 1) Let TC be the <transaction characteristics>.
- 2) TC shall contain at most one <isolation level>, at most one <transaction access mode>, and at most one <diagnostics size>.
- 3) If TC does not contain an <isolation level>, then ISOLATION LEVEL SERIALIZABLE is implicit.
- 4) If <transaction access mode> is READ WRITE, then the <level of isolation> shall not be READ UNCOMMITTED.
- 5) If TC does not contain a <transaction access mode>, then  
Case:
  - a) If <isolation level> contains READ UNCOMMITTED, then READ ONLY is implicit.
  - b) Otherwise, READ WRITE is implicit.

- 6) The declared type of <number of conditions> shall be exact numeric with scale 0 (zero).
- 7) If *TC* does not contain a <diagnostics size>, then DIAGNOSTICS SIZE *n* is implicit, where *n* is an implementation-dependent value not less than 1 (one).

## **Access Rules**

*None.*

## **General Rules**

*None.*

## **Conformance Rules**

- 1) Without Feature F111, “Isolation levels other than SERIALIZABLE”, conforming SQL language shall not contain an <isolation level> that contains a <level of isolation> other than SERIALIZABLE.
- 2) Without Feature F121, “Basic diagnostics management”, conforming SQL language shall not contain a <diagnostics size>.



## 17.4 <set constraints mode statement>

### Function

If an SQL-transaction is currently active, then set the constraint mode for that SQL-transaction in the current SQL-session. If no SQL-transaction is currently active, then set the constraint mode for the next SQL-transaction in the current SQL-session for the SQL-agent.

NOTE 434 — This statement has no effect on any SQL-transactions subsequent to this SQL-transaction.

### Format

```
<set constraints mode statement> ::=
 SET CONSTRAINTS <constraint name list> { DEFERRED | IMMEDIATE }

<constraint name list> ::=
 ALL
 | <constraint name> [{ <comma> <constraint name> }...]
```

### Syntax Rules

- 1) If a <constraint name> is specified, then it shall identify a constraint.
- 2) The constraint identified by <constraint name> shall be deferrable.

### Access Rules

*None.*

### General Rules

- 1) Let *CSC* be the current SQL-session context.
- 2) If IMMEDIATE is specified, then  
Case:
  - a) If ALL is specified, then the constraint mode in *CSC* of all constraints that are deferrable is set to immediate.
  - b) Otherwise, the constraint mode in *CSC* for the constraints identified by the <constraint name>s in the <constraint name list> is set to immediate.
- 3) If DEFERRED is specified, then  
Case:
  - a) If ALL is specified, then the constraint mode in *CSC* of all constraints that are deferrable is set to deferred.
  - b) Otherwise, the constraint mode in *CSC* for the constraints identified by the <constraint name>s in the <constraint name list> is set to deferred.

## **Conformance Rules**

- 1) Without Feature F721, “Deferrable constraints”, conforming SQL language shall not contain a <set constraints mode statement>.

## 17.5 <savepoint statement>

*This Subclause is modified by Subclause 11.17, “<savepoint statement>”, in ISO/IEC 9075-10.*

### Function

Establish a savepoint.

### Format

```
<savepoint statement> ::=
 SAVEPOINT <savepoint specifier>

<savepoint specifier> ::=
 <savepoint name>
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *S* be the <savepoint name>.
- 2) If *S* identifies an existing savepoint established within the current savepoint level, then that savepoint is destroyed.
- 3) If the number of savepoints that now exist within the current SQL-transaction is equal to the implementation-defined maximum number of savepoints per SQL-transaction, then an exception condition is raised: *savepoint exception — too many*.
- 4) A savepoint is established in the current savepoint level and at the current point in the current SQL-transaction. *S* is assigned as the identifier of that savepoint.

### Conformance Rules

- 1) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <savepoint statement>.

## 17.6 <release savepoint statement>

*This Subclause is modified by Subclause 11.18, “<release savepoint statement>”, in ISO/IEC 9075-10.*

### Function

Destroy a savepoint.

### Format

```
<release savepoint statement> ::=
 RELEASE SAVEPOINT <savepoint specifier>
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *S* be the <savepoint name>.
- 2) If *S* does not identify a savepoint established in the current savepoint level, then an exception condition is raised: *savepoint exception — invalid specification*.
- 3) The savepoint identified by *S* and all savepoints established in the current savepoint level subsequent to the establishment of *S* are destroyed.

### Conformance Rules

- 1) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <release savepoint statement>.

## 17.7 <commit statement>

This Subclause is modified by [Subclause 11.19](#), “<commit statement>”, in ISO/IEC 9075-10.

### Function

Terminate the current SQL-transaction with commit.

### Format

```
<commit statement> ::=
 COMMIT [WORK] [AND [NO] CHAIN]
```

### Syntax Rules

- 1) If neither AND CHAIN nor AND NO CHAIN is specified, then AND NO CHAIN is implicit.

### Access Rules

*None.*

### General Rules

- 1) If the current SQL-transaction is part of an encompassing transaction that is controlled by an agent other than the SQL-agent, then an exception condition is raised: *invalid transaction termination*.
- 2) If an atomic execution context is active, then an exception condition is raised: *invalid transaction termination*.
- 3) For every open cursor *CR* that is not a holdable cursor in the stack of contexts of the current SQL-session, the General Rules of [Subclause 15.4](#), “Effect of closing a cursor”, are applied with *CR* as *CURSOR* and DESTROY as *DISPOSITION*.
- 4) For every temporary table in any SQL-client module associated with the current SQL-transaction that specifies the ON COMMIT DELETE option and that was updated by the current SQL-transaction, the execution of the <commit statement> is effectively preceded by the execution of a <delete statement: searched> that specifies DELETE FROM *T*, where *T* is the <table name> of that temporary table.
- 5) The effects specified in the General Rules of [Subclause 17.4](#), “<set constraints mode statement>” occur as if the statement SET CONSTRAINTS ALL IMMEDIATE were executed for each active SQL-connection.
- 6) Case:
  - a) If any constraint is not satisfied, then any changes to SQL-data or schemas that were made by the current SQL-transaction are canceled and an exception condition is raised: *transaction rollback — integrity constraint violation*.
  - b) If any other error preventing commitment of the SQL-transaction has occurred, then any changes to SQL-data or schemas that were made by the current SQL-transaction are canceled and an exception condition is raised: *transaction rollback* with an implementation-defined subclass value.

- c) Otherwise, any changes to SQL-data or schemas that were made by the current SQL-transaction are eligible to be perceived by all concurrent and subsequent SQL-transactions.
- 7) All savepoint levels are destroyed and a new savepoint level is established.
- NOTE 435 — Destroying a savepoint level destroys all existing savepoints that are established at that level.
- 8) Every valid non-holdable locator value is marked invalid.
- 9) The current SQL-transaction is terminated.
- 10) Case:
- a) If <commit statement> contains AND CHAIN, then an SQL-transaction is initiated. Any branch transactions of the SQL-transaction are initiated with the same access mode, isolation level, and condition area limit as the corresponding branch of the SQL-transaction just terminated.
  - b) Otherwise:
    - i) The current access mode, current isolation level, and current condition area limit of the current SQL-session context are set to the access mode, isolation level, and condition area limit, respectively, of the enduring transaction characteristics of the current SQL-session.
    - ii) For every constraint *C*, the constraint mode of *C* in the current SQL-session context is set to the initial constraint mode included in the constraint descriptor for *C*.
- 11) The prepared statement of every held cursor remains in existence. It is implementation-defined whether or not any other prepared statement is deallocated.

## **Conformance Rules**

- 1) Without Feature T261, “Chained transactions”, conforming SQL language shall not contain a <commit statement> that immediately contains CHAIN.

## 17.8 <rollback statement>

This Subclause is modified by *Subclause 11.20*, “<rollback statement>”, in ISO/IEC 9075-10.

### Function

Terminate the current SQL-transaction with rollback, or rollback all actions affecting SQL-data and/or schemas since the establishment of a savepoint.

### Format

```
<rollback statement> ::=
 ROLLBACK [WORK] [AND [NO] CHAIN] [<savepoint clause>]

<savepoint clause> ::=
 TO SAVEPOINT <savepoint specifier>
```

### Syntax Rules

- 1) If AND CHAIN is specified, then <savepoint clause> shall not be specified.
- 2) If neither AND CHAIN nor AND NO CHAIN is specified, then AND NO CHAIN is implicit.

### Access Rules

*None.*

### General Rules

- 1) If the current SQL-transaction is part of an encompassing transaction that is controlled by an agent other than the SQL-agent and the <rollback statement> is not being implicitly executed, then an exception condition is raised: *invalid transaction termination*.
- 2) If a <savepoint clause> is not specified, then:
  - a) If an atomic execution context is active, then an exception condition is raised: *invalid transaction termination*.
  - b) All changes to SQL-data or schemas that were made by the current SQL-transaction are canceled.
  - c) All savepoint levels are destroyed and a new savepoint level is established.  
NOTE 436 — Destroying a savepoint level destroys all existing savepoints that are established at that level.
  - d) Every valid locator is marked invalid.
  - e) For every open cursor *CR* in the stack of contexts of the current SQL-session, the General Rules of *Subclause 15.4*, “Effect of closing a cursor”, are applied with *CR* as *CURSOR* and DESTROY as *DISPOSITION*.
  - f) It is implementation-defined whether or not any prepared statement is deallocated.

- g) The current SQL-transaction is terminated.
- h) Case:
  - i) If <rollback statement> contains AND CHAIN, then an SQL-transaction is initiated. Any branch transactions of the SQL-transaction are initiated with the same access mode, isolation level, and condition area limit as the corresponding branch of the SQL-transaction just terminated.
  - ii) Otherwise:
    - 1) The current access mode, current isolation level, and current condition area limit of the current SQL-session context are set to the access mode, isolation level, and condition area limit, respectively, of the enduring transaction characteristics of the current SQL-session.
    - 2) For every constraint *C*, the constraint mode of *C* in the current SQL-session context is set to the initial constraint mode included in the constraint descriptor for *C*.
- 3) If a <savepoint clause> is specified, then:
  - a) Let *S* be the <savepoint name>.
  - b) If *S* does not specify a savepoint established within the current savepoint level, then an exception condition is raised: *savepoint exception — invalid specification*.
  - c) If an atomic execution context is active, and *S* specifies a savepoint established before the beginning of the most recent atomic execution context, then an exception condition is raised: *savepoint exception — invalid specification*.
  - d) All changes to SQL-data or schemas that were made by the current SQL-transaction subsequent to the establishment of *S* are canceled.
  - e) All savepoints established by the current SQL-transaction subsequent to the establishment of *S* are destroyed.

NOTE 437 — Destroying a savepoint level destroys all existing savepoints that are established at that level.
  - f) Every valid locator that was generated in the current SQL-transaction subsequent to the establishment of *S* is marked invalid.
  - g) For every open cursor *CR* that is not a holdable cursor in the stack of contexts of the current SQL-session, the General Rules of Subclause 15.4, “Effect of closing a cursor”, are applied with *CR* as *CURSOR* and DESTROY as *DISPOSITION*.
  - h) The status of any open cursors in the stack of contexts of the current SQL-session that were opened by the current SQL-transaction before the establishment of *S* is implementation-defined.
  - i) It is implementation-defined whether or not any prepared statement that was prepared before the establishment of *S* is deallocated.
  - j) It is implementation-defined whether or not any prepared statement that was prepared subsequent to the establishment of *S* is deallocated.

NOTE 438 — The current SQL-transaction is not terminated, and there is no other effect on the SQL-data or schemas.

## Conformance Rules

- 1) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <savepoint clause>.



- 2) Without Feature T261, “Chained transactions”, conforming SQL language shall not contain a <rollback statement> that immediately contains CHAIN.

*(Blank page)*

## 18 Connection management

### 18.1 <connect statement>

#### Function

Establish an SQL-session.

#### Format

```
<connect statement> ::=
 CONNECT TO <connection target>

<connection target> ::=
 <SQL-server name> [AS <connection name>] [USER <connection user name>]
 | DEFAULT
```

#### Syntax Rules

- 1) If <connection user name> is not specified, then an implementation-defined <connection user name> for the SQL-connection is implicit.

#### Access Rules

*None.*

#### General Rules

- 1) If a <connect statement> is executed after the first transaction-initiating SQL-statement executed by the current SQL-transaction and the SQL-implementation does not support transactions that affect more than one SQL-server, then an exception condition is raised: *feature not supported — multiple server transactions*.
- 2) If <connection user name> is specified, then let *S* be <connection user name> and let *V* be the character string that is the value of

```
TRIM (BOTH ' ' FROM S)
```

- 3) If *V* does not conform to the Format and Syntax Rules of a <user identifier>, then an exception condition is raised: *invalid authorization specification*.
- 4) If the SQL-client module that contains the <externally-invoked procedure> that contains the <connect statement> specifies a <module authorization identifier>, then whether or not <connection user name> shall be identical to that <module authorization identifier> is implementation-defined, as are any other

restrictions on the value of <connection user name>. Otherwise, any restrictions on the value of <connection user name> are implementation-defined.

- 5) If the value of <connection user name> does not conform to the implementation-defined restrictions, then an exception condition is raised: *invalid authorization specification*.
- 6) If <connection name> was specified, then let *CV* be <simple value specification> immediately contained in <connection name>. If neither DEFAULT nor <connection name> were specified, then let *CV* be <SQL-server name>. Let *CN* be the result of

```
TRIM (BOTH ' ' FROM CV)
```

If *CN* does not conform to the Format and Syntax Rules of an <identifier>, then an exception condition is raised: *invalid connection name*.

- 7) If an SQL-connection with name *CN* has already been established by the current SQL-agent and has not been disconnected, or if DEFAULT is specified and a default SQL-connection has already been established by the current SQL-agent and has not been disconnected, then an exception condition is raised: *connection exception — connection name in use*.
- 8) Case:
  - a) If DEFAULT is specified, then the default SQL-session is initiated and associated with the default SQL-server. The method by which the default SQL-server is determined is implementation-defined.
  - b) Otherwise, an SQL-session is initiated and associated with the SQL-server identified by <SQL-server name>. The method by which <SQL-server name> is used to determine the appropriate SQL-server is implementation-defined.
- 9) If the <connect statement> successfully initiates an SQL-session, then:
  - a) The current SQL-connection *CC* and current SQL-session, if any, become a dormant SQL-connection and a dormant SQL-session, respectively. The SQL-session context for *CC* is preserved and is not affected in any way by operations performed over the initiated SQL-connection.

NOTE 439 — The SQL-session context is defined in Subclause 4.37, “SQL-sessions”.
  - b) The SQL-session initiated by the <connect statement> becomes the current SQL-session and the SQL-connection established to that SQL-session becomes the current SQL-connection.
  - c) The access mode, isolation level, and condition area limit of the enduring transaction characteristics of the current SQL-session are read-write, SERIALIZABLE, and an implementation-dependent value not less than 1 (one), respectively.

NOTE 440 — If the <connect statement> fails to initiate an SQL-session, then the current SQL-connection and current SQL-session, if any, remain unchanged.
- 10) If the SQL-client cannot establish the SQL-connection, then an exception condition is raised: *connection exception — SQL-client unable to establish SQL-connection*.
- 11) If the SQL-server rejects the establishment of the SQL-connection, then an exception condition is raised: *connection exception — SQL-server rejected establishment of SQL-connection*.
- 12) The SQL-server for the subsequent execution of <externally-invoked procedure>s in any SQL-client modules associated with the SQL-agent is set to the SQL-server identified by <SQL-server name>.
- 13) The current SQL-session context of the current SQL-session is initialized as follows:

- a) The authorization stack is set to a single cell containing the user identifier <connection user name>.
  - b) The current access mode, current isolation level, and current condition area limit are set to the access mode, isolation level, and condition area limit, respectively, of the enduring transaction characteristics of the current SQL-session.
- 14) A new savepoint level is established.

## Conformance Rules

- 1) Without Feature F771, “Connection management”, conforming SQL language shall not contain a <connect statement>.

## 18.2 <set connection statement>

### Function

Select an SQL-connection from the available SQL-connections.

### Format

```
<set connection statement> ::=
 SET CONNECTION <connection object>

<connection object> ::=
 DEFAULT
 | <connection name>
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) If a <set connection statement> is executed after the first transaction-initiating SQL-statement executed by the current SQL-transaction and the SQL-implementation does not support transactions that affect more than one SQL-server, then an exception condition is raised: *feature not supported — multiple server transactions*.
- 2) Case:
  - a) If DEFAULT is specified and there is no default SQL-connection that is current or dormant for the current SQL-agent, then an exception condition is raised: *connection exception — connection does not exist*.
  - b) Otherwise, if <connection name> does not identify an SQL-session that is current or dormant for the current SQL-agent, then an exception condition is raised: *connection exception — connection does not exist*.
- 3) If the SQL-connection identified by <connection object> cannot be selected, then an exception condition is raised: *connection exception — connection failure*.
- 4) The current SQL-connection and current SQL-session become a dormant SQL-connection and a dormant SQL-session, respectively. The SQL-session context information is preserved and is not affected in any way by operations performed over the selected SQL-connection.

NOTE 441 — The SQL-session context information is defined in [Subclause 4.37](#), “SQL-sessions”.

- 5) The SQL-connection identified by <connection object> becomes the current SQL-connection and the SQL-session associated with that SQL-connection becomes the current SQL-session. All SQL-session context information is restored to the same state as at the time the SQL-connection became dormant.

NOTE 442 — The SQL-session context information is defined in [Subclause 4.37, “SQL-sessions”](#).

- 6) The SQL-server for the subsequent execution of <externally-invoked procedure>s in any SQL-client modules associated with the SQL-agent are set to that of the current SQL-connection.

## Conformance Rules

- 1) Without Feature F771, “Connection management”, conforming SQL language shall not contain a <set connection statement>.

## 18.3 <disconnect statement>

### Function

Terminate an SQL-connection.

### Format

```
<disconnect statement> ::=
 DISCONNECT <disconnect object>

<disconnect object> ::=
 <connection object>
 | ALL
 | CURRENT
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) If <connection name> is specified and <connection name> does not identify an SQL-connection that is current or dormant for the current SQL-agent, then an exception condition is raised: *connection exception — connection does not exist*.
- 2) If DEFAULT is specified and there is no default SQL-connection that is current or dormant for the current SQL-agent, then an exception condition is raised: *connection exception — connection does not exist*.
- 3) If CURRENT is specified and there is no current SQL-connection for the current SQL-agent, then an exception condition is raised: *connection exception — connection does not exist*.
- 4) Let *C* be the current SQL-connection.
- 5) Let *L* be a list of SQL-connections. If a <connection name> is specified, then *L* is that SQL-connection. If CURRENT is specified, then *L* is the current SQL-connection. If ALL is specified, then *L* is a list representing every SQL-connection that is current or dormant for the current SQL-agent, in an implementation-dependent order. If DEFAULT is specified, then *L* is the default SQL-connection.
- 6) If any SQL-connection in *L* is active, then an exception condition is raised: *invalid transaction state — active SQL-transaction*.
- 7) For every SQL-connection *CI* in *L*, treating the SQL-session *SI* identified by *CI* as the current SQL-session, all of the actions that are required after the last call of a <externally-invoked procedure> by an SQL-agent, except for the execution of a <rollback statement> or a <commit statement>, are performed. *CI* is terminated, regardless of any exception condition that might occur during the disconnection process.



NOTE 443 — See the General Rules of Subclause 13.1, “<SQL-client module definition>”, for the actions to be performed after the last call of a <externally-invoked procedure> by an SQL-agent.

- 8) If any error is detected during execution of a <disconnect statement>, then a completion condition is raised: *warning* — *disconnect error*.
- 9) If *C* is contained in *L*, then there is no current SQL-connection following the execution of the <disconnect statement>. Otherwise, *C* remains the current SQL-connection.

## Conformance Rules

- 1) Without Feature F771, “Connection management”, conforming SQL language shall not contain a <disconnect statement>.

*(Blank page)*

## 19 Session management

*This Clause is modified by Clause 16, “Session management”, in ISO/IEC 9075-9.*

*This Clause is modified by Clause 16, “Session management”, in ISO/IEC 9075-14.*

### 19.1 <set session characteristics statement>

#### Function

Set one or more characteristics for the current SQL-session.

#### Format

```
<set session characteristics statement> ::=
 SET SESSION CHARACTERISTICS AS <session characteristic list>

<session characteristic list> ::=
 <session characteristic> [{ <comma> <session characteristic> }...]

<session characteristic> ::=
 <session transaction characteristics>

<session transaction characteristics> ::=
 TRANSACTION <transaction mode> [{ <comma> <transaction mode> }...]
```

#### Syntax Rules

- 1) <session transaction characteristics> shall contain at most one <isolation level>, at most one <transaction access mode>, and at most one <diagnostics size>.

#### Access Rules

*None.*

#### General Rules

- 1) Let *SCL* be the <session transaction characteristics>. Let *ESC* be the enduring session characteristics of the current SQL-session.
- 2) If *SCL* contains an <isolation level> *IL*, then the isolation level of *ESC* is set to the <level of isolation> contained in *IL*.

**19.1 <set session characteristics statement>**

- 3) If *SCL* contains a <transaction access mode> *AM*, then the access mode of *ESC* is set to read-only or read-write, according to whether *AM* contains READ ONLY or READ WRITE, respectively.
- 4) If *SCL* contains a <diagnostics size> *DS*, then the condition area limit of *ESC* is set to the <number of conditions> contained in *DS*.

## **Conformance Rules**

- 1) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set session characteristics statement>.

## 19.2 <set session user identifier statement>

### Function

Set the SQL-session user identifier and the current user identifier of the current SQL-session context.

### Format

```
<set session user identifier statement> ::=
 SET SESSION AUTHORIZATION <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) If a <set session user identifier statement> is executed and an SQL-transaction is currently active, then an exception condition is raised: *invalid transaction state — active SQL-transaction*.
- 2) Let *S* be <value specification> and let *V* be the character string that is the value of  

```
TRIM (BOTH ' ' FROM S)
```
- 3) If *V* does not conform to the Format and Syntax Rules of an <authorization identifier>, then an exception condition is raised: *invalid authorization specification*.
- 4) If *V* is not equal to the current value of the SQL-session user identifier of the current SQL-session context, then then restrictions on the permissible values for *V* are implementation-defined.
- 5) If the current user identifier and the current role name are restricted from setting the user identifier to *V*, then an exception condition is raised: *invalid authorization specification*.
- 6) The SQL-session user identifier of the current SQL-session context is set to *V*.
- 7) The current user identifier is set to *V*.
- 8) The SQL-session role name and the current role name are removed.

### Conformance Rules

- 1) Without Feature F321, “User authorization”, conforming SQL language shall not contain a <set session user identifier statement>.

## 19.3 <set role statement>

### Function

Set the SQL-session role name and the current role name for the current SQL-session context.

### Format

```
<set role statement> ::=
 SET ROLE <role specification>

<role specification> ::=
 <value specification>
 | NONE
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) If a <set role statement> is executed and an SQL-transaction is currently active, then an exception condition is raised: *invalid transaction state — active SQL-transaction*.
- 2) If there is no current user identifier, then an exception condition is raised: *invalid role specification*.
- 3) If <role specification> contains a <value specification>, then:
  - a) Let *S* be <value specification> and let *V* be the character string that is the value of  

```
TRIM (BOTH ' ' FROM S)
```
  - b) If *V* does not conform to the Format and Syntax Rules of a <role name>, then an exception condition is raised: *invalid role specification*.
  - c) If no role authorization descriptor exists that indicates that the role identified by *V* has been granted to either the current user identifier or to PUBLIC, then an exception condition is raised: *invalid role specification*.
  - d) The SQL-session role name and the current role name are set to *V*.
- 4) If NONE is specified, then the current role name is removed.

### Conformance Rules

- 1) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <set role statement>.

## 19.4 <set local time zone statement>

### Function

Set the current default time zone displacement for the current SQL-session.

### Format

```
<set local time zone statement> ::=
 SET TIME ZONE <set time zone value>

<set time zone value> ::=
 <interval value expression>
 | LOCAL
```

### Syntax Rules

- 1) The declared type of the <interval value expression> immediately contained in the <set time zone value> shall be INTERVAL HOUR TO MINUTE.

### Access Rules

*None.*

### General Rules

- 1) Case:
  - a) If LOCAL is specified, then the current default time zone displacement of the current SQL-session is set to the original time zone displacement of the current SQL-session.
  - b) Otherwise,  
Case:
    - i) If the value of the <interval value expression> is not the null value and is between INTERVAL -'14:00' and INTERVAL +'14:00', then the current default time zone displacement of the current SQL-session is set to the value of the <interval value expression>.
    - ii) Otherwise, an exception condition is raised: *data exception — invalid time zone displacement value*.

### Conformance Rules

- 1) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <set local time zone statement>.

## 19.5 <set catalog statement>

### Function

Set the default catalog name for unqualified <schema name>s in <preparable statement>s that are prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> and in <direct SQL statement>s that are invoked directly.

### Format

```
<set catalog statement> ::=
 SET <catalog name characteristic>

<catalog name characteristic> ::=
 CATALOG <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) Let *S* be <value specification> and let *V* be the character string that is the value of  

```
TRIM (BOTH ' ' FROM S)
```
- 2) If *V* does not conform to the Format and Syntax Rules of a <catalog name>, then an exception condition is raised: *invalid catalog name*.
- 3) The default catalog name of the current SQL-session is set to *V*.

### Conformance Rules

- 1) Without Feature F651, “Catalog name qualifiers”, conforming SQL language shall not contain a <set catalog statement>.
- 2) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set catalog statement>.



## 19.6 <set schema statement>

### Function

Set the default schema name for unqualified <schema qualified name>s in <preparable statement>s that are prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> and in <direct SQL statement>s that are invoked directly.

### Format

```
<set schema statement> ::=
 SET <schema name characteristic>

<schema name characteristic> ::=
 SCHEMA <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) Let *S* be <value specification> and let *V* be the character string that is the value of  

```
TRIM (BOTH ' ' FROM S)
```
- 2) If *V* does not conform to the Format and Syntax Rules of a <schema name>, then an exception condition is raised: *invalid schema name*.
- 3) Case:
  - a) If *V* conforms to the Format and Syntax Rules for a <schema name> that contains a <catalog name>, then let *X* be the <catalog name> part and let *Y* be the <unqualified schema name> part of *V*. The following statement is implicitly executed:  

```
SET CATALOG 'X'
```

  
and the <set schema statement> is effectively replaced by:  

```
SET SCHEMA 'Y'
```
  - b) Otherwise, the default unqualified schema name of the current SQL-session is set to *V*.

## **Conformance Rules**

- 1) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set schema statement>.

## 19.7 <set names statement>

### Function

Set the default character set name for <character string literal>s in <preparable statement>s that are prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> and in <direct SQL statement>s that are invoked directly.

### Format

```
<set names statement> ::=
 SET <character set name characteristic>

<character set name characteristic> ::=
 NAMES <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) Let *S* be <value specification> and let *V* be the character string that is the value of  

```
TRIM (BOTH ' ' FROM S)
```
- 2) If *V* does not conform to the Format and Syntax Rules of a <character set name>, then an exception condition is raised: *invalid character set name*.
- 3) The default character set name of the current SQL-session is set to *V*.

### Conformance Rules

- 1) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <set names statement>.
- 2) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set names statement>.

## 19.8 <set path statement>

### Function

Set the SQL-path used to determine the subject routine of <routine invocation>s with unqualified <routine name>s in <preparable statement>s that are prepared in the current SQL-session by an <execute immediate statement> or a <prepare statement> and in <direct SQL statement>s, respectively, that are invoked directly. The SQL-path remains the current SQL-path of the SQL-session until another SQL-path is successfully set.

### Format

```
<set path statement> ::=
 SET <SQL-path characteristic>

<SQL-path characteristic> ::=
 PATH <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) Let *S* be <value specification> and let *V* be the character string that is the value of

```
TRIM (BOTH ' ' FROM S)
```

- a) If *V* does not conform to the Format and Syntax Rules of a <schema name list>, then an exception condition is raised: *invalid schema name list specification*.
- b) The SQL-path of the current SQL-session is set to *V*.

NOTE 444 — A <set path statement> that is executed between a <prepare statement> and an <execute statement> has no effect on the prepared statement.

### Conformance Rules

- 1) Without Feature S071, “SQL paths in function and type name resolution”, Conforming SQL language shall not contain a <set path statement>.

## 19.9 <set transform group statement>

### Function

Set the group name that identifies the group of transform functions for mapping values of user-defined types to predefined data types.

### Format

```
<set transform group statement> ::=
 SET <transform group characteristic>

<transform group characteristic> ::=
 DEFAULT TRANSFORM GROUP <value specification>
 | TRANSFORM GROUP FOR TYPE <path-resolved user-defined type name> <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.
- 2) If <path-resolved user-defined type name> is specified, then let *UDT* be the user-defined type identified by that <path-resolved user-defined type name>.

### Access Rules

*None.*

### General Rules

- 1) Let *S* be <value specification> and let *V* be the character string that is the value of

```
TRIM (BOTH ' ' FROM S)
```

- a) If *V* does not conform to the Format and Syntax Rules of a <group name>, then an exception condition is raised: *invalid transform group name specification*.
- b) Case:
  - i) If <path-resolved user-defined type name> is specified, then the transform group name corresponding to all subtypes of *UDT* for the current SQL-session is set to *V*.
  - ii) Otherwise, the default transform group name for the current SQL-session is set to *V*.

NOTE 445 — A <set transform group statement> that is executed after a <prepare statement> has no effect on the prepared statement.

### Conformance Rules

- 1) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <set transform group statement>.

## 19.10 <set session collation statement>

### Function

Set the SQL-session collation of the SQL-session for one or more character sets. An SQL-session collation remains effective until another SQL-session collation for the same character set is successfully set.

### Format

```
<set session collation statement> ::=
 SET COLLATION <collation specification> [FOR <character set specification list>]
 | SET NO COLLATION [FOR <character set specification list>]

<collation specification> ::=
 <value specification>
```

### Syntax Rules

- 1) The declared type of the <value specification> shall be a character string type.

### Access Rules

*None.*

### General Rules

- 1) Let *S* be <value specification> and let *V* be the character string that is the value of

```
TRIM (BOTH ' ' FROM S)
```

- a) If *V* does not conform to the Format and Syntax Rules of a <collation name>, then an exception condition is raised: *invalid collation name*.
- b) Let *CO* be the collation identified by the <collation name> contained in *V*.

Case:

- i) If <character set specification list> is specified, then

Case:

- 1) If the collation specified by *CO* is not applicable to any character set identified by a <character set specification>, then an exception condition is raised: *invalid collation name*.
- 2) Otherwise, for each character set specified, the SQL-session collation for that character set in the current SQL-session is set to *CO*.

- ii) Otherwise, the character sets for which the SQL-session collations are set to *CO* are implementation-defined.

- 2) If SET NO COLLATION is specified, then

Case:

- a) If <character set specification list> is specified, then, for each character set specified, the SQL-session collation for that character set in the current SQL-session is set to *none*.
- b) Otherwise, the SQL-session collation for every character set in the current SQL-session is set to *none*.

## Conformance Rules

- 1) Without Feature F693, “SQL-session and client module collations”, conforming SQL language shall not contain a <set session collation statement>.

*(Blank page)*



## 20 Dynamic SQL

### 20.1 Description of SQL descriptor areas

*This Subclause is modified by Subclause 17.1, “Description of SQL descriptor areas”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 17.1, “Description of SQL descriptor areas”, in ISO/IEC 9075-14.*

#### Function

Specify the identifiers, data types, and codes used in SQL item descriptor areas.

#### Syntax Rules

- 1) An SQL item descriptor area comprises the items specified in Table 24, “Data types of <key word>s used in SQL item descriptor areas”.
- 2) An SQL descriptor area comprises the items specified in Table 23, “Data types of <key word>s used in the header of SQL descriptor areas”, and one or more occurrences of an SQL item descriptor area.
- 3) Given an SQL item descriptor area *IDA* in which the value of LEVEL is *N*, the *immediately subordinate descriptor areas* of *IDA* are those SQL item descriptor areas in which the value of LEVEL is *N+1* and whose position in the SQL descriptor area follows that of *IDA* and precedes that of any SQL item descriptor area in which the value of LEVEL is less than *N+1*.

The *subordinate descriptor areas* of *IDA* are those SQL item descriptor areas that are immediately subordinate descriptor areas of *IDA* or that are subordinate descriptor areas of an SQL item descriptor area that is immediately subordinate to *IDA*.

- 4) Given a data type *DT* and its descriptor *DE*, the *immediately subordinate descriptors* of *DE* are defined to be

Case:

- a) If *DT* is a row type, then the field descriptors of the fields of *DT*. The *i*-th immediately subordinate descriptor is the descriptor of the *i*-th field of *DT*.
- b) If *DT* is a collection type, then the descriptor of the associated element type of *DT*.

The *subordinate descriptors* of *DE* are those descriptors that are immediately subordinate descriptors of *DE* or that are subordinate descriptors of a descriptor that is immediately subordinate to *DE*.

- 5) Given a descriptor *DE*, let *SDE<sub>j</sub>* represent its *j*-th immediately subordinate descriptor. There is an implied ordering of the subordinate descriptors of *DE*, such that:
  - a) *SDE<sub>1</sub>* is in the first ordinal position.

## 20.1 Description of SQL descriptor areas

- b) The ordinal position of  $SDE_{j+1}$  is  $K+NS+1$ , where  $K$  is the ordinal position of  $SDE_j$  and  $NS$  is the number of subordinate descriptors of  $SDE_j$ . The implicitly ordered subordinate descriptors of  $SDE_j$  occupy contiguous ordinal positions starting at position  $K+1$ .
- 6) An item descriptor area  $IDA$  is *valid* if and only if TYPE indicates a code defined in Table 25, “Codes used for SQL data types in Dynamic SQL”, and one of the following is true.

Case:

- a) TYPE indicates CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT, LENGTH is a valid length value for TYPE, and CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA, and CHARACTER\_SET\_NAME are the fully qualified name of a character set that is valid for TYPE.
- b) TYPE indicates CHARACTER LARGE OBJECT LOCATOR.
- c) TYPE indicates BINARY, BINARY VARYING, or BINARY LARGE OBJECT and LENGTH is a valid length value for the TYPE.
- d) TYPE indicates BINARY LARGE OBJECT LOCATOR.
- e) TYPE indicates NUMERIC and PRECISION and SCALE are valid precision and scale values for the NUMERIC data type.
- f) TYPE indicates DECIMAL and PRECISION and SCALE are valid precision and scale values for the DECIMAL data type.
- g) TYPE indicates SMALLINT, INTEGER, BIGINT, REAL, or DOUBLE PRECISION.
- h) TYPE indicates FLOAT and PRECISION is a valid precision value for the FLOAT data type.
- i) TYPE indicates BOOLEAN.
- j) TYPE indicates a <datetime type>, DATETIME\_INTERVAL\_CODE is a code specified in Table 26, “Codes associated with datetime data types in Dynamic SQL”, and PRECISION is a valid value for the <time precision> or <timestamp precision> of the indicated datetime data type.
- k) TYPE indicates an <interval type>, DATETIME\_INTERVAL\_CODE is a code specified in Table 27, “Codes used for <interval qualifier>s in Dynamic SQL”, and DATETIME\_INTERVAL\_PRECISION and PRECISION are valid values for <interval leading field precision> and <interval fractional seconds precision> for an <interval qualifier>.
- l) TYPE indicates USER-DEFINED TYPE LOCATOR and USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME are the fully qualified name of a valid user-defined type.
- m) TYPE indicates REF, LENGTH is the length in octets for the REF type, and USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME are a valid qualified user-defined type name, and SCOPE\_CATALOG, SCOPE\_SCHEMA, and SCOPE\_NAME are a valid qualified table name.
- n) TYPE indicates ROW, the value  $N$  of DEGREE is a valid value for the degree of a row type, there are exactly  $N$  immediately subordinate descriptor areas of  $IDA$  and those SQL item descriptor areas are valid.

## 20.1 Description of SQL descriptor areas

- o) TYPE indicates ARRAY or ARRAY LOCATOR, the value of CARDINALITY is a valid value for the cardinality of an array, there is exactly one immediately subordinate descriptor area of *IDA*, and that SQL item descriptor area is valid.
  - p) TYPE indicates MULTISET or MULTISET LOCATOR, there is exactly one immediately subordinate descriptor area of *IDA*, and that SQL item descriptor area is valid.
  - q) 09 TYPE indicates an implementation-defined data type.
- 7) The declared type *T* of a <simple value specification> or a <simple target specification> *SVT* is said to *match* the data type specified by a valid item descriptor area *IDA* if and only if one of the following conditions is true.

Case:

- a) TYPE indicates CHARACTER and *T* is specified by CHARACTER(*L*), where *L* is the value of LENGTH and the <character set specification> formed by the values of CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA, and CHARACTER\_SET\_NAME identifies the character set of *SVT*.
- b) Either TYPE indicates CHARACTER VARYING and *T* is specified by CHARACTER VARYING(*L*) or type indicates CHARACTER LARGE OBJECT and *T* is specified by CHARACTER LARGE OBJECT(*L*), where the <character set specification> formed by the values of CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA, and CHARACTER\_SET\_NAME identifies the character set of *SVT* and

Case:

- i) *SVT* is a <simple value specification> and *L* is the value of LENGTH.
- ii) *SVT* is a <simple target specification> and *L* is not less than the value of LENGTH.
- c) TYPE indicates CHARACTER LARGE OBJECT LOCATOR and *T* is specified by CHARACTER LARGE OBJECT LOCATOR.
- d) TYPE indicates BINARY and *T* is specified by BINARY(*L*), where *L* is the value of LENGTH.
- e) Either TYPE indicates BINARY VARYING and *T* is specified by BINARY VARYING or TYPE indicates BINARY LARGE OBJECT and *T* is specified by BINARY LARGE OBJECT(*L*), and

Case:

- i) *STV* is a <simple value specification> and *L* is the value of LENGTH.
- ii) *STV* is a <simple target specification> and *L* is not less than the value of LENGTH.
- f) TYPE indicates BINARY LARGE OBJECT LOCATOR and *T* is specified by BINARY LARGE OBJECT LOCATOR.
- g) TYPE indicates NUMERIC and *T* is specified by NUMERIC(*P*,*S*), where *P* is the value of PRECISION and *S* is the value of SCALE.
- h) TYPE indicates DECIMAL and *T* is specified by DECIMAL(*P*,*S*), where *P* is the value of PRECISION and *S* is the value of SCALE.
- i) TYPE indicates SMALLINT and *T* is specified by SMALLINT.
- j) TYPE indicates INTEGER and *T* is specified by INTEGER.

## 20.1 Description of SQL descriptor areas

- k) TYPE indicates BIGINT and  $T$  is specified by BIGINT.
- l) TYPE indicates FLOAT and  $T$  is specified by  $\text{FLOAT}(P)$ , where  $P$  is the value of PRECISION.
- m) TYPE indicates REAL and  $T$  is specified by REAL.
- n) TYPE indicates DOUBLE PRECISION and  $T$  is specified by DOUBLE PRECISION.
- o) TYPE indicates BOOLEAN and  $T$  is specified by BOOLEAN.
- p) TYPE indicates USER-DEFINED TYPE LOCATOR and  $T$  is specified by USER-DEFINED TYPE LOCATOR, where the values of USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME are the fully qualified name of the associated user-defined type of SVT.
- q) TYPE indicates REF and  $T$  is specified by REF, where the <user-defined type name> formed by the values of USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME identifies the row type of SVT, and SCOPE\_CATALOG, SCOPE\_SCHEMA, and SCOPE\_NAME identify the scope of the reference type.
- r) TYPE indicates ROW, and  $T$  is a row type with degree  $D$ , where  $D$  is the value of DEGREE, and the data type of the  $i$ -th field of SVT matches the data type specified by the  $i$ -th immediately subordinate descriptor area of IDA.
- s) TYPE indicates ARRAY and  $T$  is an array type with maximum cardinality  $C$  and the data type of the element type of  $T$  matches the immediately subordinate descriptor area of IDA, and

Case:

- i) SVT is a <simple value specification> and  $C$  is the value of CARDINALITY.
- ii) SVT is a <simple target specification> and  $C$  is not less than the value of CARDINALITY.
- t) TYPE indicates ARRAY LOCATOR and  $T$  is an array locator type whose associated array type has maximum cardinality  $C$  and the data type of the element type of the associated array type of  $T$  matches the immediately subordinate descriptor area of IDA, and

Case:

- i) SVT is a <simple value specification> and  $C$  is the value of CARDINALITY.
  - ii) SVT is a <simple target specification> and  $C$  is not less than the value of CARDINALITY.
  - u) TYPE indicates MULTISET and  $T$  is a multiset type and the data type of the element type of  $T$  matches the immediately subordinate descriptor area of IDA.
  - v) TYPE indicates MULTISET LOCATOR and  $T$  is a multiset locator type and the data type of the element type of  $T$  matches the immediately subordinate descriptor area of IDA.
  - w) TYPE indicates a data type from Table 25, “Codes used for SQL data types in Dynamic SQL”, other than an implementation-defined data type and  $T$  satisfies the implementation-defined rules for matching that data type.
  - x) 09 TYPE indicates an implementation-defined data type and  $T$  satisfies the implementation-defined rules for matching that data type.
- 8) A data type  $DT$  is said to be *represented* by an SQL item descriptor area if a <simple value specification> of type  $DT$  matches the SQL item descriptor area.

**Table 23 — Data types of <key word>s used in the header of SQL descriptor areas**

<b>&lt;key word&gt;</b>	<b>Data Type</b>
COUNT	exact numeric with scale 0 (zero)
DYNAMIC_FUNCTION	character string with character set SQL_IDENTIFIER and length not less than 128 characters
DYNAMIC_FUNCTION_CODE	exact numeric with scale 0 (zero)
KEY_TYPE	exact numeric with scale 0 (zero)
TOP_LEVEL_COUNT	exact numeric with scale 0 (zero)

**Table 24 — Data types of <key word>s used in SQL item descriptor areas**

<b>&lt;key word&gt;</b>	<b>Data Type</b>
CARDINALITY	exact numeric with scale 0 (zero)
CHARACTER_SET_CATALOG	character string with character set SQL_IDENTIFIER and length not less than 128 characters
CHARACTER_SET_NAME	character string with character set SQL_IDENTIFIER and length not less than 128 characters
CHARACTER_SET_SCHEMA	character string with character set SQL_IDENTIFIER and length not less than 128 characters
COLLATION_CATALOG	character string with character set SQL_IDENTIFIER and length not less than 128 characters
COLLATION_NAME	character string with character set SQL_IDENTIFIER and length not less than 128 characters
COLLATION_SCHEMA	character string with character set SQL_IDENTIFIER and length not less than 128 characters
DATA	matches the data type represented by the SQL item descriptor area
DATETIME_INTERVAL_CODE	exact numeric with scale 0 (zero)
DATETIME_INTERVAL_PRECISION	exact numeric with scale 0 (zero)
DEGREE	exact numeric with scale 0 (zero)

<key word>	Data Type
INDICATOR	exact numeric with scale 0 (zero)
KEY_MEMBER	exact numeric with scale 0 (zero)
LENGTH	exact numeric with scale 0 (zero)
LEVEL	exact numeric with scale 0 (zero)
NAME	character string with character set SQL_IDENTIFIER and length not less than 128 characters
NULLABLE	exact numeric with scale 0 (zero)
OCTET_LENGTH	exact numeric with scale 0 (zero)
PARAMETER_MODE	exact numeric with scale 0 (zero)
PARAMETER_ORDINAL_POSITION	exact numeric with scale 0 (zero)
PARAMETER_SPECIFIC_CATALOG	character string with character set SQL_IDENTIFIER and length not less than 128 characters
PARAMETER_SPECIFIC_NAME	character string with character set SQL_IDENTIFIER and length not less than 128 characters
PARAMETER_SPECIFIC_SCHEMA	character string with character set SQL_IDENTIFIER and length not less than 128 characters
PRECISION	exact numeric with scale 0 (zero)
RETURNED_CARDINALITY	exact numeric with scale 0 (zero)
RETURNED_LENGTH	exact numeric with scale 0 (zero)
RETURNED_OCTET_LENGTH	exact numeric with scale 0 (zero)
SCALE	exact numeric with scale 0 (zero)
SCOPE_CATALOG	character string with character set SQL_IDENTIFIER and length not less than 128 characters
SCOPE_NAME	character string with character set SQL_IDENTIFIER and length not less than 128 characters
SCOPE_SCHEMA	character string with character set SQL_IDENTIFIER and length not less than 128 characters
TYPE	exact numeric with scale 0 (zero)
UNNAMED	exact numeric with scale 0 (zero)

<key word>	Data Type
USER_DEFINED_TYPE_CATALOG	character string with character set SQL_IDENTIFIER and length not less than 128 characters
USER_DEFINED_TYPE_NAME	character string with character set SQL_IDENTIFIER and length not less than 128 characters
USER_DEFINED_TYPE_SCHEMA	character string with character set SQL_IDENTIFIER and length not less than 128 characters
USER_DEFINED_TYPE_CODE	exact numeric with scale 0 (zero)

NOTE 446 — “Matches” and “represented by”, as applied to the relationship between a data type and an SQL item descriptor area are defined in the Syntax Rules of this Subclause.

## Access Rules

*None.*

## General Rules

- 1) [0914](#) Table 25, “Codes used for SQL data types in Dynamic SQL”, specifies the codes associated with the SQL data types.

[0914](#) Table 25 — Codes used for SQL data types in Dynamic SQL

Data Type	Code
Implementation-defined data types	< 0 (zero)
ARRAY	50
ARRAY LOCATOR	51
BIGINT	25
BINARY	60
BINARY VARYING	61
BINARY LARGE OBJECT	30
BINARY LARGE OBJECT LOCATOR	31
BOOLEAN	16
CHARACTER	1 (one)

Data Type	Code
CHARACTER VARYING	12
CHARACTER LARGE OBJECT	40
CHARACTER LARGE OBJECT LOCATOR	41
DATE, TIME WITHOUT TIME ZONE, TIME WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, or TIMESTAMP WITH TIME ZONE	9
DECIMAL	3
DOUBLE PRECISION	8
FLOAT	6
INTEGER	4
INTERVAL	10
MULTISET	55
MULTISET LOCATOR	56
NUMERIC	2
REAL	7
SMALLINT	5
USER-DEFINED TYPE LOCATOR	18
ROW TYPE	19
REF	20
User-defined types	17

- 2) Table 26, “Codes associated with datetime data types in Dynamic SQL”, specifies the codes associated with the datetime data types.

**Table 26 — Codes associated with datetime data types in Dynamic SQL**

Datetime Data Type	Code
DATE	1 (one)



Datetime Data Type	Code
TIME WITH TIME ZONE	4
TIME WITHOUT TIME ZONE	2
TIMESTAMP WITH TIME ZONE	5
TIMESTAMP WITHOUT TIME ZONE	3

- 3) Table 27, “Codes used for <interval qualifier>s in Dynamic SQL”, specifies the codes associated with <interval qualifier>s for interval data types.

**Table 27 — Codes used for <interval qualifier>s in Dynamic SQL**

Interval Qualifier	Code
DAY	3
DAY TO HOUR	8
DAY TO MINUTE	9
DAY TO SECOND	10
HOUR	4
HOUR TO MINUTE	11
HOUR TO SECOND	12
MINUTE	5
MINUTE TO SECOND	13
MONTH	2
SECOND	6
YEAR	1 (one)
YEAR TO MONTH	7

- 4) The value of DYNAMIC\_FUNCTION is a character string that identifies the type of the prepared or executed SQL-statement. Table 32, “SQL-statement codes”, specifies the identifier of the SQL-statements.
- 5) The value of DYNAMIC\_FUNCTION\_CODE is a number that identifies the type of the prepared or executed SQL-statement. Table 32, “SQL-statement codes”, specifies the identifier of the SQL-statements.

**20.1 Description of SQL descriptor areas**

- 6) Table 28, “Codes used for input/output SQL parameter modes in Dynamic SQL”, specifies the codes used for the PARAMETER\_MODE item descriptor field when describing a <call statement>.

**Table 28 — Codes used for input/output SQL parameter modes in Dynamic SQL**

Parameter mode	Code
PARAMETER_MODE_IN	1 (one)
PARAMETER_MODE_INOUT	2
PARAMETER_MODE_OUT	4

- 7) Table 29, “Codes associated with user-defined types in Dynamic SQL”, specifies the codes associated with user-defined types.

**Table 29 — Codes associated with user-defined types in Dynamic SQL**

User-Defined Type	Code
DISTINCT	1 (one)
STRUCTURED	2

**Conformance Rules**

*None.*

## 20.2 <allocate descriptor statement>

### Function

Allocate an SQL descriptor area.

### Format

```
<allocate descriptor statement> ::=
 ALLOCATE [SQL] DESCRIPTOR <descriptor name> [WITH MAX <occurrences>]

<occurrences> ::=
 <simple value specification>
```

### Syntax Rules

- 1) The declared type of <occurrences> shall be exact numeric with scale 0 (zero).
- 2) If WITH MAX <occurrences> is not specified, then an implementation-defined default value for <occurrences> that is greater than 0 (zero) is implicit.

### Access Rules

*None.*

### General Rules

- 1) Case:
  - a) If a <descriptor name> is an <extended descriptor name>, then let *S* be the <simple value specification> that is immediately contained in <extended descriptor name> and let *V* be the character string that is the result of  
  
TRIM ( BOTH ' ' FROM *S* )  
  
Case:
    - i) If *V* does not conform to the Format and Syntax Rules of an <identifier>, then an exception condition is raised: *invalid SQL descriptor name*.
    - ii) Otherwise, let *DN* be the <extended descriptor name>. The value of *DN* is *V*.
  - b) Otherwise, let *DN* be the <non-extended descriptor name>.
- 2) If <occurrences> is less than 1 (one) or is greater than an implementation-defined maximum value, then an exception condition is raised: *dynamic SQL error — invalid descriptor index*. The maximum number of SQL descriptor areas that can be allocated at one time is implementation-defined.
- 3) Case:

- a) If *DN* identifies an SQL descriptor area, then an exception condition is raised: *invalid SQL descriptor name*.
- b) Otherwise, an SQL descriptor area is created that is identified by *DN*. The SQL descriptor area will have at least <occurrences> number of SQL item descriptor areas. The value of LEVEL in each of the item descriptor areas is set to 0 (zero). The value of every other field in the SQL descriptor area is implementation-dependent.

## Conformance Rules

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain an <occurrences> that is not a <literal>.
- 2) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <allocate descriptor statement>.

## 20.3 <deallocate descriptor statement>

### Function

Deallocate an SQL descriptor area.

### Format

```
<deallocate descriptor statement> ::=
 DEALLOCATE [SQL] DESCRIPTOR <descriptor name>
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Case:
  - a) If <descriptor name> does not identify an SQL descriptor area, then an exception condition is raised:  
*invalid SQL descriptor name.*
  - b) Otherwise, the SQL descriptor area identified by <descriptor name> is destroyed.

### Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <deallocate descriptor statement>.

## 20.4 <get descriptor statement>

### Function

Get information from an SQL descriptor area.

### Format

```

<get descriptor statement> ::=
 GET [SQL] DESCRIPTOR <descriptor name> <get descriptor information>

<get descriptor information> ::=
 <get header information> [{ <comma> <get header information> }...]
 | VALUE <item number> <get item information>
 [{ <comma> <get item information> }...]

<get header information> ::=
 <simple target specification 1> <equals operator> <header item name>

<header item name> ::=
 COUNT
 | KEY_TYPE
 | DYNAMIC_FUNCTION
 | DYNAMIC_FUNCTION_CODE
 | TOP_LEVEL_COUNT

<get item information> ::=
 <simple target specification 2> <equals operator> <descriptor item name>

<item number> ::=
 <simple value specification>

<simple target specification 1> ::=
 <simple target specification>

<simple target specification 2> ::=
 <simple target specification>

<descriptor item name> ::=
 CARDINALITY
 | CHARACTER_SET_CATALOG
 | CHARACTER_SET_NAME
 | CHARACTER_SET_SCHEMA
 | COLLATION_CATALOG
 | COLLATION_NAME
 | COLLATION_SCHEMA
 | DATA
 | DATETIME_INTERVAL_CODE
 | DATETIME_INTERVAL_PRECISION
 | DEGREE
 | INDICATOR
 | KEY_MEMBER
 | LENGTH
 | LEVEL
 | NAME
 | NULLABLE

```

```
OCTET_LENGTH
PARAMETER_MODE
PARAMETER_ORDINAL_POSITION
PARAMETER_SPECIFIC_CATALOG
PARAMETER_SPECIFIC_NAME
PARAMETER_SPECIFIC_SCHEMA
PRECISION
RETURNED_CARDINALITY
RETURNED_LENGTH
RETURNED_OCTET_LENGTH
SCALE
SCOPE_CATALOG
SCOPE_NAME
SCOPE_SCHEMA
TYPE
UNNAMED
USER_DEFINED_TYPE_CATALOG
USER_DEFINED_TYPE_NAME
USER_DEFINED_TYPE_SCHEMA
USER_DEFINED_TYPE_CODE
```

## Syntax Rules

- 1) The declared type of <item number> shall be exact numeric with scale 0 (zero).
- 2) For each <get header information>, the declared type of <simple target specification 1> shall be that shown in the Data Type column of the row in Table 23, “Data types of <key word>s used in the header of SQL descriptor areas”, whose <key word> column value is equivalent to <header item name>.
- 3) For each <get item information>, the declared type of <simple target specification 2> shall be that shown in the Data Type column of the row in Table 24, “Data types of <key word>s used in SQL item descriptor areas”, whose <key word> column value is equivalent to <descriptor item name>.

## Access Rules

*None.*

## General Rules

- 1) If <descriptor name> does not identify an SQL descriptor area, then an exception condition is raised: *invalid SQL descriptor name*.
- 2) If the <item number> specified in a <get descriptor statement> is greater than the value of <occurrences> specified when the <descriptor name> was allocated or less than 1 (one), then an exception condition is raised: *dynamic SQL error — invalid descriptor index*.
- 3) If the <item number> specified in a <get descriptor statement> is greater than the value of COUNT, then a completion condition is raised: *no data*.
- 4) If the declared type of the <simple target specification> associated with the keyword DATA does not match the data type represented by the item descriptor area, then an exception condition is raised: *data exception — error in assignment*.

NOTE 447 — “Match” and “represented by” are defined in the Syntax Rules of Subclause 20.1, “Description of SQL descriptor areas”.

- 5) Let  $i$  be the value of the <item number> contained in <get descriptor information>. Let  $IDA$  be the  $i$ -th item descriptor area. If a <get item information> specifies DATA, then:
  - a) If  $IDA$  is subordinate to an item descriptor area whose TYPE field indicates ARRAY, ARRAY LOCATOR, MULTISET, or MULTISET LOCATOR, then an exception condition is raised: *dynamic SQL error — undefined DATA value*.
  - b) If the value of TYPE in  $IDA$  indicates ROW, then an exception condition is raised: *dynamic SQL error — undefined DATA value*.
  - c) If the value of INDICATOR is negative and no <get item information> specifies INDICATOR, then an exception condition is raised: *data exception — null value, no indicator parameter*.
- 6) If an exception condition is raised in a <get descriptor statement>, then the values of all targets specified by <simple target specification 1> and <simple target specification 2> are implementation-dependent.
- 7) A <get descriptor statement> retrieves values from the SQL descriptor area and item specified by <descriptor name>. For each item, the value that is retrieved is the one established by the most recently executed <allocate descriptor statement>, <set descriptor statement>, or <describe statement> that references the specified SQL descriptor area and item. The value retrieved by a <get descriptor statement> for any field whose value is undefined is implementation-dependent.

Case:

- a) If <get descriptor information> contains one or more <get header information>s, then for each <get header information> specified, the value of <simple target specification 1> is set to the value  $V$  in the SQL descriptor area of the field identified by the <header item name> by applying the General Rules of Subclause 9.2, “Store assignment”, to <simple target specification 1> and  $V$  as TARGET and VALUE, respectively.
- b) If <get descriptor information> contains one or more <get item information>s, then:
  - i) Let  $i$  be the value of the <item number> contained in the <get descriptor information>.
  - ii) For each <get item information> specified, the value of <simple target specification 2> is set to the value  $V$  in the  $i$ -th SQL item descriptor area of the field identified by the <descriptor item name> by applying the General Rules of Subclause 9.2, “Store assignment”, to <simple target specification 2> and  $V$  as TARGET and VALUE, respectively.

## Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <get descriptor statement>.
- 2) Without Feature T301, “Functional dependencies”, conforming SQL language shall not contain a <descriptor item name> that contains KEY\_MEMBER.



## 20.5 <set descriptor statement>

### Function

Set information in an SQL descriptor area.

### Format

```
<set descriptor statement> ::=
 SET [SQL] DESCRIPTOR <descriptor name> <set descriptor information>

<set descriptor information> ::=
 <set header information> [{ <comma> <set header information> }...]
 | VALUE <item number> <set item information>
 [{ <comma> <set item information> }...]

<set header information> ::=
 <header item name> <equals operator> <simple value specification 1>

<set item information> ::=
 <descriptor item name> <equals operator> <simple value specification 2>

<simple value specification 1> ::=
 <simple value specification>

<simple value specification 2> ::=
 <simple value specification>
```

### Syntax Rules

- 1) For each <set header information>, <header item name> shall not be KEY\_TYPE, TOP\_LEVEL\_COUNT, DYNAMIC\_FUNCTION, or DYNAMIC\_FUNCTION\_CODE, and the declared type of <simple value specification 1> shall be that in the Data Type column of the row of [Table 23, “Data types of <key word>s used in the header of SQL descriptor areas”](#), whose <key word> column value is equivalent to <header item name>.
- 2) For each <set item information>, the value of <descriptor item name> shall not be RETURNED\_LENGTH, RETURNED\_OCTET\_LENGTH, RETURNED\_CARDINALITY, OCTET\_LENGTH, NULLABLE, KEY\_MEMBER, COLLATION\_CATALOG, COLLATION\_SCHEMA, COLLATION\_NAME, NAME, UNNAMED, PARAMETER\_MODE, PARAMETER\_ORDINAL\_POSITION, PARAMETER\_SPECIFIC\_CATALOG, PARAMETER\_SPECIFIC\_SCHEMA, PARAMETER\_SPECIFIC\_NAME, or USER\_DEFINED\_TYPE\_CODE. Other alternatives for <descriptor item name> shall not be specified more than once in a <set descriptor statement>. The declared type of <simple value specification 2> shall be that shown in the Data Type column of the row in [Table 24, “Data types of <key word>s used in SQL item descriptor areas”](#), whose <key word> column value is equivalent to <descriptor item name>.
- 3) If the <descriptor item name> specifies DATA, then <simple value specification 2> shall not be a <literal>.

### Access Rules

*None.*

## General Rules

- 1) If <descriptor name> does not identify an SQL descriptor area, then an exception condition is raised: *invalid SQL descriptor name*.
- 2) If the <item number> specified in a <set descriptor statement> is greater than the value of <occurrences> specified when the <descriptor name> was allocated or less than 1 (one), then an exception condition is raised: *dynamic SQL error — invalid descriptor index*.
- 3) When more than one value is set in a single <set descriptor statement>, the values are effectively assigned in the following order: LEVEL, TYPE, DATETIME\_INTERVAL\_CODE, DATETIME\_INTERVAL\_PRECISION, PRECISION, SCALE, CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA, CHARACTER\_SET\_NAME, USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, USER\_DEFINED\_TYPE\_NAME, SCOPE\_CATALOG, SCOPE\_SCHEMA, SCOPE\_NAME, LENGTH, INDICATOR, DEGREE, CARDINALITY, and DATA.

When any value other than DATA is set, the value of DATA becomes undefined.

- 4) For every <set item information> specified, let *DIN* be the <descriptor item name>, let *V* be the value of the <simple value specification 2>, let *N* be the value of <item number>, and let *IDA* be the *N*-th item descriptor area.

Case:

- a) If *DIN* is DATA, then:
  - i) If *IDA* is subordinate to an item descriptor area whose TYPE field indicates ARRAY, ARRAY LOCATOR, MULTiset, or MULTiset LOCATOR, then an exception condition is raised: *dynamic SQL error — invalid DATA target*.
  - ii) If TYPE in *IDA* indicates ROW, then an exception condition is raised: *dynamic SQL error — invalid DATA target*.
  - iii) If the most specific type of *V* does not match the data type specified by the item descriptor area, then an exception condition is raised: *data exception — error in assignment*.

NOTE 448 — “Match” is defined in the Syntax Rules of [Subclause 20.1](#), “Description of SQL descriptor areas”.

- iv) The value of DATA in *IDA* is set to *V*.
- b) If *DIN* is LEVEL, then:
  - i) If *N* is 1 (one) and *V* is not 0 (zero), then an exception condition is raised: *dynamic SQL error — invalid LEVEL value*.
  - ii) If *N* is greater than 1 (one), then let *PIDA* be *IDA*'s immediately preceding item descriptor area and let *K* be its LEVEL value.
    - 1) If  $V = K + 1$  and TYPE in *PIDA* does not indicate ROW, ARRAY, ARRAY LOCATOR, MULTiset, MULTiset LOCATOR, then an exception condition is raised: *dynamic SQL error — invalid LEVEL value*.
    - 2) If  $V > K + 1$ , then an exception condition is raised: *dynamic SQL error — invalid LEVEL value*.
    - 3) If  $V < K + 1$ , then let *OIDA<sub>i</sub>* be the *i*-th item descriptor area to which *PIDA* is subordinate and whose TYPE field indicates ROW, let *NS<sub>i</sub>* be the number of immediately subordinate

descriptor areas of  $OIDA_i$  between  $OIDA_i$  and  $IDA$  and let  $D_i$  be the value of DEGREE in  $OIDA_i$ .

- A) For each  $OIDA_i$  whose LEVEL value is greater than  $V$ , if  $D_i$  is not equal to  $NS_i$ , then an exception condition is raised: *dynamic SQL error — invalid LEVEL value*.
- B) If  $K$  is not 0 (zero), then let  $OIDA_j$  be the  $OIDA_i$  whose LEVEL value is  $K$ . If there exists no such  $OIDA_j$  or  $D_j$  is not greater than  $NS_j$ , then an exception condition is raised: *dynamic SQL error — invalid LEVEL value*.

iii) The value of LEVEL in  $IDA$  is set to  $V$ .

c) If  $DIN$  is TYPE, then:

- i) The value of TYPE in  $IDA$  is set to  $V$ .
- ii) The value of all fields other than TYPE and LEVEL in  $IDA$  are set to implementation-dependent values.
- iii) Case:
  - 1) If  $V$  indicates CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT, then CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA and CHARACTER\_SET\_NAME in  $IDA$  are set to the values for the default character set name for the SQL-session and LENGTH in  $IDA$  is set to 1 (one).
  - 2) If  $V$  indicates CHARACTER LARGE OBJECT LOCATOR, then LENGTH in  $IDA$  is set to 1 (one).
  - 3) If  $V$  indicates BINARY, BINARY VARYING, or BINARY LARGE OBJECT, then LENGTH in  $IDA$  is set to 1 (one).
  - 4) If  $V$  indicates BINARY LARGE OBJECT LOCATOR, then LENGTH in  $IDA$  is set to 1 (one).
  - 5) If  $V$  indicates DATETIME, then PRECISION in  $IDA$  is set to 0 (zero).
  - 6) If  $V$  indicates INTERVAL, then DATETIME\_INTERVAL\_PRECISION in  $IDA$  is set to 2.
  - 7) If  $V$  indicates NUMERIC or DECIMAL, then SCALE in  $IDA$  is set to 0 (zero) and PRECISION in  $IDA$  is set to the implementation-defined default value for the precision of NUMERIC or DECIMAL data types, respectively.
  - 8) If  $V$  indicates FLOAT, then PRECISION in  $IDA$  is set to the implementation-defined default value for the precision of the FLOAT data type.

d) If  $DIN$  is DATETIME\_INTERVAL\_CODE, then

Case:

- i) If TYPE in  $IDA$  indicates DATETIME, then

Case:

- 1) If  $V$  indicates DATE, TIME, or TIME WITH TIME ZONE, then PRECISION in  $IDA$  is set to 0 (zero) and DATETIME\_INTERVAL\_CODE in  $IDA$  is set to  $V$ .

- 2) If *V* indicates **TIMESTAMP** or **TIMESTAMP WITH TIME ZONE**, then **PRECISION** in *IDA* is set to 6 and **DATETIME\_INTERVAL\_CODE** in *IDA* is set to *V*.
  - 3) Otherwise, an exception condition is raised: *dynamic SQL error — invalid DATE-TIME\_INTERVAL\_CODE*.
- ii) If **TYPE** in *IDA* indicates **INTERVAL**, then
- Case:
- 1) If *V* indicates **DAY TO SECOND**, **hour to second**, **minute to second**, or **second**, then **PRECISION** in *IDA* is set to 6, **DATETIME\_INTERVAL\_PRECISION** in *IDA* is set to 2 and **DATETIME\_INTERVAL\_CODE** in *IDA* is set to *V*.
  - 2) If *V* indicates **YEAR**, **MONTH**, **DAY**, **hour**, **minute**, **year to month**, **day to hour**, **day to minute**, or **hour to minute**, then **PRECISION** in *IDA* is set to 0 (zero), **DATETIME\_INTERVAL\_PRECISION** in *IDA* is set to 2 and **DATETIME\_INTERVAL\_CODE** in *IDA* is set to *V*.
  - 3) Otherwise, an exception condition is raised: *dynamic SQL error — invalid DATE-TIME\_INTERVAL\_CODE*.
- iii) Otherwise, an exception condition is raised: *dynamic SQL error — invalid DATETIME\_INTERVAL\_CODE*.
- e) Otherwise, the value of *DIN* in *IDA* is set to *V* by applying the General Rules of [Subclause 9.2, “Store assignment”](#), to the field of *IDA* identified by *DIN* and *V* as *TARGET* and *VALUE*, respectively.
- 5) For each <set header information> specified, the value of the field identified by <header item name> is set to the value *V* of <simple value specification 1> by applying the General Rules of [Subclause 9.2, “Store assignment”](#), to the field identified by the <header item name> and *V* as *TARGET* and *VALUE*, respectively.
  - 6) If an exception condition is raised in a <set descriptor statement>, then the values of all elements of the item descriptor area specified in the <set descriptor statement> are implementation-dependent.
  - 7) Restrictions on changing **TYPE**, **LENGTH**, **PRECISION**, **SCALE**, **DATETIME\_INTERVAL\_CODE**, **DATETIME\_INTERVAL\_PRECISION**, **CHARACTER\_SET\_CATALOG**, **CHARACTER\_SET\_SCHEMA**, **CHARACTER\_SET\_NAME**, **USER\_DEFINED\_TYPE\_CATALOG**, **USER\_DEFINED\_TYPE\_SCHEMA**, **USER\_DEFINED\_TYPE\_NAME**, **SCOPE\_CATALOG**, **SCOPE\_SCHEMA**, and **SCOPE\_NAME** values resulting from the execution of a <describe statement> before execution of an <execute statement>, <dynamic open statement>, or <dynamic fetch statement> are implementation-defined.

## Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <set descriptor statement>.

## 20.6 <prepare statement>

*This Subclause is modified by Subclause 15.1, “<prepare statement>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 17.2, “<prepare statement>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 17.4, “<prepare statement>”, in ISO/IEC 9075-14.*

### Function

Prepare a statement for execution.

### Format

```
<prepare statement> ::=
 PREPARE <SQL statement name> [<attributes specification>]
 FROM <SQL statement variable>

<attributes specification> ::=
 ATTRIBUTES <attributes variable>

<attributes variable> ::=
 <simple value specification>

<SQL statement variable> ::=
 <simple value specification>

<preparable statement> ::=
 <preparable SQL data statement>
 | <preparable SQL schema statement>
 | <preparable SQL transaction statement>
 | <preparable SQL control statement>
 | <preparable SQL session statement>
 | <preparable implementation-defined statement>

<preparable SQL data statement> ::=
 <delete statement: searched>
 | <dynamic single row select statement>
 | <insert statement>
 | <dynamic select statement>
 | <update statement: searched>
 | <truncate table statement>
 | <merge statement>
 | <preparable dynamic delete statement: positioned>
 | <preparable dynamic update statement: positioned>
 | <hold locator statement>
 | <free locator statement>

<preparable SQL schema statement> ::=
 <SQL schema statement>

<preparable SQL transaction statement> ::=
 <SQL transaction statement>

<preparable SQL control statement> ::=
 <SQL control statement>
```

```
<preparable SQL session statement> ::=
 <SQL session statement>

<dynamic select statement> ::=
 <cursor specification>

<preparable implementation-defined statement> ::=
 !! See the Syntax Rules.
```

## Syntax Rules

- 1) The <simple value specification> of <SQL statement variable> shall not be a <literal>.
- 2) The declared types of each of <SQL statement variable> and <attributes variable> shall be character string.
- 3) The Format and Syntax Rules for <preparable implementation-defined statement> are implementation-defined.
- 4) A <preparable SQL control statement> shall not contain an <SQL procedure statement> that is not a <preparable statement>, nor shall it contain a <dynamic single row select statement> or a <dynamic select statement>.

## Access Rules

*None.*

## General Rules

- 1) 09 Let *P* be the contents of the <SQL statement variable>. If *P* is an <SQL control statement>, then let *PS* be an <SQL procedure statement> contained in *P*.
- 2) Two subfields *SF1* and *SF2* of row types *RT1* and *RT2* are *corresponding subfields* if *SF1* and *SF2* are positionally corresponding fields of *RT1* and *RT2*, respectively, or if *SF1* and *SF2* are positionally corresponding fields of *RT1SF1* and *RT2SF2* and *RT1SF1* and *RT2SF2* are the declared types of corresponding subfields of *RT1* and *RT2*, respectively.
- 3) Let *DTGN* be the default transform group name and let *TFL* be the list of {user-defined type name — transform group name} pairs used to identify the group of transform functions for every user-defined type that is referenced in *P*. *DTGN* and *TFL* are not affected by the execution of a <set transform group statement> after *P* is prepared.
- 4) Let *DPV* be a <value expression> that is either a <dynamic parameter specification> or a <dynamic parameter specification> immediately contained in any number of <left paren> <right paren> pairs. Initially, the declared type of such a <value expression> is, by definition, undefined. A data type is *undefined* if it is neither a data type defined in this standard nor a data type defined by the implementation.
- 5) 04 Let *MP* be the implementation-defined maximum value of <precision> for the NUMERIC data type. Let *ML* be the implementation-defined maximum length of varying-length character strings. For each <value expression> *DP* in *P* or *PS* that meets the criteria for *DPV* let *DT* denote its declared type. The syntactic substitutions specified in Subclause 14.15, “<set clause list>”, shall not be applied until the data types of <dynamic parameter specification>s are determined by this General Rule.
  - a) Case:

- i) If *DP* is immediately followed by an <interval qualifier> *IQ*, then *DT* is INTERVAL *IQ*.
- ii) If *DP* is the <numeric value expression> simply contained in an <array element reference>, then *DT* is NUMERIC (*MP*, 0).
- iii) If *DP* is the <string value expression> simply contained in a <char length expression> or an <octet length expression>, then *DT* is CHARACTER VARYING(*ML*) with an implementation-defined character set.
- iv) If *DP* is either the <numeric value expression dividend> *X1* or the <numeric value expression divisor> *X2* simply contained in a <modulus expression>, then if *DP* is *X1* (*X2*), then *DT* is the declared type of *X2* (*X1*).
- v) If *DP* is either *X1* or *X2* in a <position expression> of the form “POSITION <left paren> *X1* IN *X2* <right paren>”, and if *DP* is *X1* (*X2*), then

Case:

- 1) If the declared type of *X2* (*X1*) is CHARACTER or CHARACTER VARYING with character set *CS*, then *DT* is CHARACTER VARYING (*ML*) with character set *CS*.
- 2) Otherwise, *DT* is the declared type of *X2* (*X1*).
- vi) If *DP* is either *X2* or *X3* in a <string value function> of the form “SUBSTRING <left paren> *X1* FROM *X2* FOR *X3* <right paren>” or “SUBSTRING <left paren> *X1* FROM *X2* <right paren>”, then *DT* is NUMERIC (*MP*, 0).
- vii) If *DP* is either *X1*, *X2*, or *X3* in a <string value function> of the form “SUBSTRING (*X1* SIMILAR *X2* ESCAPE *X3* )”, then
  - 1) Case:
    - A) If the declared type of *X1* is CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT, then let *CS* be the character set of *X1*.
    - B) If the declared type of *X2* is CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT, then let *CS* be the character set of *X1*.
    - C) If the declared type of *X3* is CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT, then let *CS* be the character set of *X1*.
    - D) Otherwise, the character set of *CS* is undefined.
  - 2) If *CS* is defined, then:
    - A) If *DP* is *X1* or *X2*, then *DT* is CHARACTER VARYING(*ML*) with character set *CS*.
    - B) If *DP* is *X3*, then *DT* is CHARACTER(1) with character set *CS*.
- viii) If *DP* is any of *X1*, *X2*, *X3*, or *X4* in a <string value function> of the form “OVERLAY <left paren> *X1* PLACING *X2* FROM *X3* FOR *X4* <right paren>” or “OVERLAY <left paren> *X1* PLACING *X2* FROM *X3* <right paren>”, then

Case:

- 1) If *DP* is *X1* (*X2*), then

Case:

- A) If the declared type of  $X_2$  ( $X_1$ ) is CHARACTER or CHARACTER VARYING with character set  $CS$ ,  $DT$  is CHARACTER VARYING ( $ML$ ) with character set  $CS$ .
- B) Otherwise,  $DT$  is the declared type of  $X_2$  ( $X_1$ ).
- 2) Otherwise,  $DT$  is NUMERIC ( $MP$ , 0).
- ix) If  $DP$  is either  $X_1$  or  $X_2$  in a <value expression> of the form “ $X_1$  <concatenation operator>  $X_2$ ” and  $DP$  is  $X_1$  ( $X_2$ ), then  
Case:
  - 1) If the declared type of  $X_2$  ( $X_1$ ) is CHARACTER or CHARACTER VARYING with character set  $CS$ , then  $DT$  is CHARACTER VARYING ( $ML$ ) with character set  $CS$ .
  - 2) Otherwise,  $DT$  is the declared type of  $X_2$  ( $X_1$ ).
- x) If  $DP$  is either  $X_1$  or  $X_2$  in a <value expression> of the form “ $X_1$  <asterisk>  $X_2$ ” or “ $X_1$  <solidus>  $X_2$ ” and  $DP$  is  $X_1$  ( $X_2$ ), then  
Case:
  - 1) If  $DP$  is  $X_1$ , then  $DT$  is the declared type of  $X_2$ .
  - 2) Otherwise,  
Case:
    - A) If the declared type of  $X_1$  is an interval type, then  $DT$  is NUMERIC ( $MP$ , 0).
    - B) Otherwise,  $DT$  is the declared type of  $X_2$  ( $X_1$ ).
- xi) If  $DP$  is either  $X_1$  or  $X_2$  in a <value expression> of the form “ $X_1$  <plus sign>  $X_2$ ” or “ $X_1$  <minus sign>  $X_2$ ”, then  
Case:
  - 1) If  $DP$  is  $X_1$  in an expression of the form “ $X_1$  <minus sign>  $X_2$ ”, then  $DT$  is the declared type of  $X_2$ .
  - 2) Otherwise, if  $DP$  is  $X_1$  ( $X_2$ ), then  
Case:
    - A) If the declared type of  $X_2$  ( $X_1$ ) is date, then  $DT$  is INTERVAL YEAR ( $PR$ ) TO MONTH, where  $PR$  is the implementation-defined maximum <interval leading field precision>.
    - B) If the declared type of  $X_2$  ( $X_1$ ) is time or timestamp, then  $DT$  is INTERVAL DAY ( $PR$ ) TO SECOND( $FR$ ), where  $PR$  and  $FR$  are the implementation-defined maximum <interval leading field precision> and maximum <interval fractional seconds precision>, respectively.
    - C) Otherwise,  $DT$  is the declared type of  $X_2$  ( $X_1$ ).
- xii) If  $DP$  is the <value expression primary> simply contained in a <boolean primary>, then  $DT$  is BOOLEAN.
- xiii) If  $DP$  is an <array element> simply contained in an <array element list>  $AEL$  or  $DP$  represents the value of a subfield  $SF$  of the declared type of an <array element> simply contained in an



<array element list> *AEL*, then let *ET* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the <array element>s simply contained in *AEL*.

Case:

- 1) If *DP* is an <array element> of *AEL*, then *DT* is *ET*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *ET* that corresponds to *SF*.
- xiv) If *DP* is a <multiset element> simply contained in a <multiset element list> *MEL* or *DP* represents the value of a subfield *SF* of the declared type of a <multiset element> simply contained in a <multiset element list> *MEL*, then let *ET* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the <multiset element>s simply contained in *MEL*.

Case:

- 1) If *DP* is a <multiset element> of *MEL*, then *DT* is *ET*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *ET* that corresponds to *SF*.
- xv) If *DP* is the <cast operand> simply contained in a <cast specification> *CS* or *DP* represents the value of a subfield *SF* of the declared type of the <cast operand> simply contained in a <cast specification> *CS*, then let *CT* be the simply contained <cast target> of *CS*.

- 1) Let *RT* be a data type determined as follows.

Case:

- A) If *CT* immediately contains ARRAY or MULTISET, then *RT* is undefined.
- B) If *CT* immediately contains <data type>, then *RT* is that data type.
- C) If *CT* simply contains <domain name> *D*, then *RT* is the declared type of the domain identified by *D*.

- 2) Case:

- A) If *DP* is the <cast operand> of *CS*, *DT* is *RT*.
- B) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.

- xvi) If *DP* is a <value expression> simply contained in a <case abbreviation> *CA* or *DP* represents the value of a subfield *SF* of the declared type of such a <value expression>, then let *RT* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the <value expression>s simply contained in *CA*.

Case:

- 1) If *DP* is a <value expression> simply contained in *CA*, then *DT* is *RT*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.
- xvii) If *DP* is a <result expression> simply contained in a <case specification> *CE* or *DP* represents the value of a subfield *SF* of the declared type of such a <result expression>, then let *RT* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the <result expression>s simply contained in *CE*.

Case:

- 1) If *DP* is a <result expression> simply contained in *CE*, then *DT* is *RT*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.
- xviii) If *DP* is a <case operand> or <when operand> simply contained in a <simple case> *CE* or *DP* represents the value of a subfield *SF* of the declared type of such a <case operand> or <when operand>, then *RT* is the result of applying the Syntax Rules of [Subclause 9.3, “Result of data type combinations”](#), to the declared types of the <case operand> and <when operand>s simply contained in *CE*.

Case:

- 1) If *DP* is a <case operand> or <when operand> simply contained in *CE*, then *DT* is *RT*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.
- xix) If *DP* is a <row value expression> or <contextually typed row value expression> simply contained in a <table value constructor> or <contextually typed table value constructor> *TVC*, or if *DP* represents the value of a subfield *SF* of the declared type of such a <row value expression> or <contextually typed row value expression>, then

Case:

- 1) Let *RT* be a data type determined as follows:

Case:

- A) If *TVC* is simply contained in a <query expression> that is simply contained in an <insert statement> *IS* or if *TVC* is immediately contained in the <insert columns and source> of an <insert statement> *IS*, then *RT* is a row type in which the declared type of the *i*-th field is the declared type of the *i*-th column in the explicit or implicit <insert column list> of *IS* and the degree of *RT* is equal to the number of columns in the explicit or implicit <insert column list> of *IS*.
- B) Otherwise, *RT* is the result of applying the Syntax Rules of [Subclause 9.3, “Result of data type combinations”](#), to the declared types of the <row value expression>s or <contextually typed row value expression>s simply contained in *TVC*.

- 2) Case:

- A) If *DP* is a <row value expression> or <contextually typed row value expression> simply contained in *TVC*, then *DT* is *RT*.
- B) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.

- xx) If *DP* is the <value expression> simply contained in an <merge insert value list> of an <merge insert specification> *MIS* of a <merge statement> or if *DP* represents the value of a subfield *SF* of the declared type of such a <value expression>, then let *RT* be the data type indicated in the column descriptor for the positionally corresponding column in the explicit or implicit <insert column list> contained in *MIS*.

Case:

- 1) If *DP* is the <value expression> simply contained in *MIS*, then *DT* is *RT*.
- 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.

- xxi) If *DP* is a <row value predicand> simply contained in a <comparison predicate>, <distinct predicate> or <between predicate> *PR* or if *DP* represents the value of a subfield *SF* of the declared type of such a <row value predicand>, then let *RT* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the <row value predicand>s simply contained in *PR*.

Case:

- 1) If *DP* is a <row value predicand> simply contained in *PR*, then *DT* is *RT*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.
- xxii) If *DP* is a <row value predicand> simply contained in a <quantified comparison predicate> or <match predicate> *PR* or *DP* represents the value of a subfield *SF* of the declared type of such a <row value predicand>, then let *RT* be the declared type of the <table subquery> simply contained in *PR*.

Case:

- 1) If *DP* is a <row value predicand> simply contained in *PR*, then *DT* is *RT*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.
- xxiii) If *DP* is a <row value predicand> simply contained in an <in predicate> *PR* or if *DP* represents the value of a subfield *SF* of the declared type of such a <row value predicand>, then let *RT* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of the <row value predicand>s simply contained in *PR* and the declared row type of the <table subquery> (if any) simply contained in *PR*.

Case:

- 1) If *DP* is a <row value predicand> simply contained in *PR*, then *DT* is *RT*.
  - 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.
- xxiv) If *DP* is the first <row value constructor element> simply contained in either <row value predicand 1> *RV1* or <row value predicand 2> *RV2* in an <overlaps predicate> *PR*, then

Case:

- 1) If both *RV1* and *RV2* simply contain a <row value constructor predicand> whose first <row value constructor element> meets the criteria for *DPV*, then *DT* is **TIMESTAMP WITH TIME ZONE**.
  - 2) Otherwise, if *DP* is simply contained in *RV1* (*RV2*), then *DT* is the declared type of the first field of *RV2* (*RV1*).
- xxv) If *DP* is simply contained in a <character like predicate>, <octet like predicate>, or <similar predicate> *PR*, then let *X1* represent the <row value predicand> immediately contained in *PR*, let *X2* represent the <character pattern>, the <octet pattern> or the <similar pattern>, and let *X3* represent the <escape character> or the <escape octet>.

Case:

- 1) If all *X1*, *X2* and *X3* meet the criteria for *DPV*, then *DT* is **CHARACTER VARYING (ML)** with an implementation-defined character set.

- 2) Otherwise, let *RT* be the result of applying the Syntax Rules of Subclause 9.3, “Result of data type combinations”, to the declared types of *X1*, *X2* and *X3*.

Case:

- A) If *RT* is CHARACTER or CHARACTER VARYING with character set *CS*, then *DT* is CHARACTER VARYING(*ML*) with character set *CS*.

- B) Otherwise, *DT* is *RT*.

- xxvi) If *DP* is the <value expression> simply contained in an <update source> of a <set clause> *SC* or if *DP* represents the value of a subfield *SF* of the declared type of such a <value expression>, then let *RT* be the declared type of the <update target> or <mutated set clause> specified in *SC*.

Case:

- 1) If *DP* is the <value expression> simply contained in *SC*, then *DT* is *RT*.
- 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.

- xxvii) 04 If *DP* is a <contextually typed row value expression> simply contained in a <multiple column assignment> *MCA* of a <set clause> *SC* or if *DP* represents the value of a subfield *SF* of the declared type of such a <contextually typed row value expression>, then let *RT* be a row type in which the declared type of the *i*-th field is the declared type of the <update target> or <mutated set clause> immediately contained in the *i*-th <set target> contained in the <set target list> of *MCA*.

Case:

- 1) If *DP* is a <contextually typed row value expression> simply contained in *MCA*, then *DT* is *RT*.
- 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.

- xxviii) If *DP* is the <value specification> immediately contained in a <catalog name characteristic>, <schema name characteristic>, <character set name characteristic>, <SQL-path characteristic>, <transform group characteristic>, <role specification> or <set session user identifier statement>, then *DT* is CHARACTER VARYING (*ML*) with an implementation-defined character set.

- xxix) If *DP* is the <interval value expression> immediately contained in a <set local time zone statement>, then *DT* is INTERVAL HOUR TO MINUTE.

- xxx) The Syntax Rules of Subclause 10.4, “<routine invocation>”, are invoked with a <routine invocation> *RI* as *ROUTINE INVOCATION*, yielding subject routine *SR*.

- xxxi) 04 If *DP* is an <SQL argument> of *RI* or if *DP* is the value of a subfield *SF* of the declared type of a <value expression> immediately contained in such an <SQL argument>, and if *DP* is the *i*-th <SQL argument> of *RI* or is contained in the *i*-th <SQL argument> of *RI*, then let *RT* denote the declared type of the *i*-th SQL parameter of *SR*.

Case:

- 1) If *DP* is the *i*-th <SQL argument> of *RI*, then *DT* is *RT*.
- 2) Otherwise, *DT* is the declared type of the subfield of *RT* that corresponds to *SF*.

- xxxii) If *DP* is contained in a <window frame preceding> or a <window frame following> contained in a <window specification> *WS*, then

Case:

- 1) If *WS* specifies ROWS, then *DT* is NUMERIC (*MP*, 0).
- 2) Otherwise, let *SDT* be the data type of the single <sort key> contained in *WS*.

Case:

- A) If *SDT* is a numeric type, then *DT* is *SDT*.
- B) If *SDT* is DATE, then *DT* is INTERVAL DAY.
- C) If *SDT* is TIME(*P*) WITHOUT TIME ZONE or TIME(*P*) WITH TIME ZONE, then *DT* is INTERVAL HOUR TO SECOND(*P*).
- D) If *SDT* is TIMESTAMP(*P*) WITHOUT TIME ZONE or TIMESTAMP(*P*) WITH TIME ZONE, then *DT* is INTERVAL DAY TO SECOND(*P*).
- E) If *SDT* is an interval type, then *DT* is *SDT*.

xxxiii) If *DP* is a <locator reference> simply contained in a <hold locator statement> or a <free locator statement>, then *DT* is INTEGER.

xxxiv) If *DP* is an <XQuery pattern>, <XQuery option flag>, <regex subject string>, or <XQuery replacement string> immediately contained in a <regex occurrences function>, <regex position expression>, <regex substring function>, <regex transliteration>, or <regex like predicate> *FUN* and if there exists an <XQuery pattern>, <XQuery option flag>, <regex subject string>, or <XQuery replacement string> immediately contained in *FUN*, then let *CS* be the character set of the <XQuery pattern>, <XQuery option flag>, <regex subject string>, or <XQuery replacement string> that is not a <dynamic parameter specification>. *DT* is CHARACTER VARYING (*ML*) CHARACTER SET *CS*.

xxxv) 14 If *DP* is a <start position>, <regex occurrence>, or <regex capture group> immediately contained in a <regex occurrences function>, <regex position expression>, <regex substring function>, or <regex transliteration>, then *DT* is NUMERIC(*MP*).

b) If *DT* is undefined, then an exception condition is raised: *syntax error or access rule violation*.

- 6) If *P* does not conform to the Format, Syntax Rules, and Access Rules of a <preparable statement>, or if *P* contains a <simple comment> then

Case:

- a) If *P* contains a <preparable dynamic cursor name> that is ambiguous, then an exception condition is raised: *ambiguous cursor name*.
- b) If *P* contains a <preparable dynamic cursor name> that is invalid, then an exception condition is raised: *invalid cursor name*.
- c) Otherwise, an exception condition is raised: *syntax error or access rule violation*.

- 7) Whether a <dynamic parameter specification> is an input argument, an output argument, or both an input and an output argument is determined as follows:

Case:

- a) If *P* is a <call statement>, then:

- i) Let *SR* be the subject routine of the <routine invocation> *RI* immediately contained in *P*. Let *n* be the number of <SQL argument>s in the <SQL argument list> immediately contained in *RI*.
  - ii) Let *A<sub>y</sub>*, 1 (one) ≤ *y* ≤ *n*, be the *y*-th <SQL argument> of the <SQL argument list> immediately contained in *RI*.
  - iii) For each <dynamic parameter specification> *D* contained in some <SQL argument> *A<sub>k</sub>*, 1 (one) ≤ *k* ≤ *n*:
    - 1) *D* is an input <dynamic parameter specification> if the <parameter mode> of the *k*-th SQL parameter of *SR* of *SR* is IN or INOUT.
    - 2) *D* is an output <dynamic parameter specification> if the <parameter mode> of the *k*-th SQL parameter of *SR* is OUT or INOUT.
  - b) Otherwise:
    - i) If a <dynamic parameter specification> is contained in a <target specification>, then it is an output <dynamic parameter specification>.
    - ii) If a <dynamic parameter specification> is contained in a <value specification>, then it is an input <dynamic parameter specification>.
- 8) If <extended statement name> is specified for the <SQL statement name>, then let *S* be <simple value specification> and let *V* be the character string that is the result of

```
TRIM (BOTH ' ' FROM S)
```

Case:

- a) If *V* does not conform to the Format and Syntax Rules of an <identifier>, then an exception condition is raised: *invalid SQL statement identifier*.
  - b) Otherwise, let *ESN* be the <extended statement name>. The value of *ESN* is *V*.
- 9) If <SQL statement name> identifies a prepared statement *PS*, then an implicit

```
DEALLOCATE PREPARE SSN
```

is executed, where *SSN* is an <SQL statement name> that identifies *PS*.

- 10) *P* is prepared for execution, resulting in a prepared statement *PRP*.

Case:

- a) If the <prepare statement> is contained in an <SQL-invoked routine> *R*, then
 

Case:

  - i) If the security characteristic of *R* is DEFINER, then the owner of *PRP* is set to the owner of *R*.
  - ii) Otherwise, *PRP* has no owner.
- b) If the <prepare statement> is contained in a triggered action, then the owner of *PRP* is set to the owner of the trigger.
- c) Otherwise,

NOTE 449 — If the <prepare statement> is in neither of the above, then it must necessarily be immediately contained in an externally-invoked procedure.

Case:

- i) If the SQL-client module that includes the <prepare statement> has a <module authorization identifier> *MAI* and FOR STATIC ONLY was not specified in the <SQL-client module definition>, then the owner of *PRP* is *MAI*.
  - ii) Otherwise, *PRP* has no owner.
- 11) If <attributes specification> is specified, then let *ATV* be the contents of the <attributes variable>. If *ATV* is not a zero-length character string and if *ATV* does not conform to the Format and Syntax Rules of Subclause 20.7, “<cursor attributes>”, then an exception condition is raised: *syntax error or access rule violation*.

## Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <prepare statement>.
- 2) Without Feature B034, “Dynamic specification of cursor attributes”, conforming SQL language shall not contain an <attributes specification>.

## 20.7 <cursor attributes>

### Function

Specify a list of cursor attributes.

### Format

```
<cursor attributes> ::=
 <cursor attribute>...

<cursor attribute> ::=
 <cursor sensitivity>
 | <cursor scrollability>
 | <cursor holdability>
 | <cursor returnability>
```

### Syntax Rules

- 1) Each of <cursor sensitivity>, <cursor scrollability>, <cursor holdability> and <cursor returnability> shall be specified at most once.

### Access Rules

*None.*

### General Rules

*None.*

### Conformance Rules

*None.*



## 20.8 <deallocate prepared statement>

This Subclause is modified by *Subclause 17.3*, “<deallocate prepared statement>”, in ISO/IEC 9075-9.

### Function

Deallocate SQL-statements that have been prepared with a <prepare statement>.

### Format

```
<deallocate prepared statement> ::=
 DEALLOCATE PREPARE <SQL statement name>
```

### Syntax Rules

- 1) If <SQL statement name> is a <statement name>, then

Case:

- a) If the <deallocate prepared statement> is contained in an <SQL-invoked routine>, then the innermost containing <SQL-invoked routine> shall contain a <prepare statement> whose <statement name> is equivalent to the <statement name> of the <deallocate prepared statement>.
- b) Otherwise, the containing <SQL-client module definition> shall contain a <prepare statement> whose <statement name> is equivalent to the <statement name> of the <deallocate prepared statement>.

### Access Rules

*None.*

### General Rules

- 1) 09 If <SQL statement name> does not identify a prepared statement, then an exception condition is raised: *invalid SQL statement name*.
- 2) If <SQL statement name> identifies a prepared statement that is the <cursor specification> of an open cursor, then an exception condition is raised: *invalid cursor state*.
- 3) The prepared statement identified by the <SQL statement name> is destroyed. The cursor instance descriptor of any declared dynamic cursor that is associated with the prepared statement is destroyed. The cursor declaration descriptor and cursor instance descriptor of any extended statement cursor that is associated with the prepared statement identified by the <SQL statement name> is destroyed. If the value of the <SQL statement name> identifies an existing prepared statement that is a <cursor specification>, then any prepared statements that reference that cursor are destroyed.

## **Conformance Rules**

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <deallocate prepared statement>.

## 20.9 <describe statement>

*This Subclause is modified by Subclause 17.4, “<describe statement>”, in ISO/IEC 9075-9.*

### Function

Obtain information about the <select list> columns or <dynamic parameter specification>s contained in a prepared statement or about the columns of the result set associated with a cursor.

### Format

```
<describe statement> ::=
 <describe input statement>
 | <describe output statement>

<describe input statement> ::=
 DESCRIBE INPUT <SQL statement name> <using descriptor> [<nesting option>]

<describe output statement> ::=
 DESCRIBE [OUTPUT] <described object> <using descriptor> [<nesting option>]

<nesting option> ::=
 WITH NESTING
 | WITHOUT NESTING

<using descriptor> ::=
 USING [SQL] DESCRIPTOR <descriptor name>

<described object> ::=
 <SQL statement name>
 | CURSOR <extended cursor name> STRUCTURE
```

### Syntax Rules

- 1) If <SQL statement name> is a <statement name>, then  
Case:
  - a) If the <describe statement> is contained in an <SQL-invoked routine>, then the innermost containing <SQL-invoked routine> shall contain a <prepare statement> whose <statement name> is equivalent to the <statement name> of the <describe statement>.
  - b) Otherwise, the containing <SQL-client module definition> shall contain a <prepare statement> whose <statement name> is equivalent to the <statement name> of the <describe statement>.
- 2) If <nesting option> is not specified, then WITHOUT NESTING is implicit.

### Access Rules

*None.*

## General Rules

- 1) 09 If <SQL statement name> is specified and does not identify a prepared statement, then an exception condition is raised: *invalid SQL statement name*.
- 2) If <extended cursor name> is specified and does not identify an extended dynamic cursor, then an exception condition is raised: *invalid cursor name*.
- 3) If <descriptor name> does not identify an SQL descriptor area, then an exception condition is raised: *invalid SQL descriptor name*.
- 4) Let  $DA$  be the descriptor area identified by <descriptor name>. Let  $N$  be the <occurrences> specified when  $DA$  was allocated.
- 5) Case:
  - a) If the statement being executed is a <describe input statement>, then a descriptor for the input <dynamic parameter specification>s for the prepared statement is stored in  $DA$ . Let  $D$  be the number of input <dynamic parameter specification>s in the prepared statement. If WITH NESTING is specified, then let  $NS_i$ ,  $1 \text{ (one)} \leq i \leq D$ , be the number of subordinate descriptors of the descriptor for the  $i$ -th input dynamic parameter; otherwise, let  $NS_i$  be 0 (zero).
  - b) If the statement being executed is a <describe output statement> and the prepared statement that is being described is a <dynamic select statement> or a <dynamic single row select statement>, then a descriptor for the <select list> columns for the prepared statement is stored in  $DA$ . Let  $T$  be the table defined by the prepared statement and let  $D$  be the degree of  $T$ . If WITH NESTING is specified, then let  $NS_i$ ,  $1 \text{ (one)} \leq i \leq D$ , be the number of subordinate descriptors of the descriptor for the  $i$ -th column of  $T$ ; otherwise, let  $NS_i$  be 0 (zero).
  - c) Otherwise, a descriptor for the output <dynamic parameter specification>s for the prepared statement is stored in  $DA$ . Let  $D$  be the number of output <dynamic parameter specification>s in the prepared statement. If WITH NESTING is specified, then let  $NS_i$ ,  $1 \text{ (one)} \leq i \leq D$ , be the number of subordinate descriptors of the descriptor for the  $i$ -th output dynamic parameter; otherwise, let  $NS_i$  be 0 (zero).
- 6)  $DA$  is set as follows:
  - a) Let  $TD$  be the value of  $D+NS_1+NS_2+\dots+NS_D$ . COUNT is set to  $TD$ .
  - b) TOP\_LEVEL\_COUNT is set to  $D$ .
  - c) DYNAMIC\_FUNCTION and DYNAMIC\_FUNCTION\_CODE are set to the identifier and code, respectively, for the prepared statement as shown in Table 32, “SQL-statement codes”. It is implementation-defined whether the identifier and code from Table 32, “SQL-statement codes”, for <dynamic select statement> or <dynamic single row select statement> is used to describe a <dynamic select statement> or <dynamic single row select statement> that has been prepared but has not yet been executed dynamically.
  - d) If the statement being executed is a <describe output statement> and the prepared statement that is being described is a <dynamic select statement> or a <dynamic single row select statement>, then  
Case:
    - i) If some subset of the columns of  $T$  is the primary key of  $T$ , then KEY\_TYPE is set to 1 (one).

- ii) If some subset of the columns of  $T$  is the preferred candidate key of  $T$ , then KEY\_TYPE is set to 2.
- iii) Otherwise, KEY\_TYPE is set to 0 (zero).

NOTE 450 — Primary keys and preferred candidate keys are defined in Subclause 4.18, “Functional dependencies”.

- e) If  $TD$  is greater than  $N$ , then a completion condition is raised: *warning — insufficient item descriptor areas*.
- f) If  $TD$  is 0 (zero) or  $TD$  is greater than  $N$ , then no item descriptor areas are set. Otherwise:
  - i) The first  $TD$  item descriptor areas are set with values from the descriptors and, optionally, subordinate descriptors for
 

Case:

    - 1) If the statement being executed is a <describe input statement>, then the input <dynamic parameter specification>s.
    - 2) If the statement being executed is a <describe output statement> and the statement being described is a <dynamic select statement> or a <dynamic single row select statement>, then the columns of  $T$ .
    - 3) Otherwise, the output <dynamic parameter specification>s.
  - ii) The descriptor for the first such column or <dynamic parameter specification> is assigned to the first item descriptor area.
  - iii) If the descriptor for the  $j$ -th column or <dynamic parameter specification> is assigned to the  $k$ -th item descriptor area, then:
    - 1) The descriptor for the  $(j+1)$ -th column or <dynamic parameter specification> is assigned to the  $(k+NS_{j+1})$ -th item descriptor area.
    - 2) If WITH NESTING is specified, then the implicitly ordered subordinate descriptors for the  $j$ -th column or <dynamic parameter specification> are assigned to contiguous item descriptor areas starting at the  $(k+1)$ -th item descriptor area.

- 7) An SQL item descriptor area, if set, consists of values for LEVEL, TYPE, NULLABLE, NAME, UNNAMED, PARAMETER\_ORDINAL\_POSITION, PARAMETER\_SPECIFIC\_CATALOG, PARAMETER\_SPECIFIC\_SCHEMA, PARAMETER\_SPECIFIC\_NAME, and other fields depending on the value of TYPE as described below. The DATA and INDICATOR fields are not relevant. Those fields and fields that are not applicable for a particular value of TYPE are set to implementation-dependent values.
  - a) If the SQL item descriptor area is set to a descriptor that is immediately subordinate to another whose LEVEL value is  $K$ , then LEVEL is set to  $K+1$ ; otherwise, LEVEL is set to 0 (zero).
  - b) TYPE is set to a code, as shown in Table 25, “Codes used for SQL data types in Dynamic SQL”, indicating the declared type of the column, <dynamic parameter specification>, or subordinate descriptor.
  - c) Case:
    - i) If the value of LEVEL is 0 (zero) and the item descriptor area describes a column, then:
      - 1) If the column is possibly nullable, then NULLABLE is set to 1 (one); otherwise, NULLABLE is set to 0 (zero).

- 2) If the column name is implementation-dependent, then NAME is set to the implementation-dependent name of the column and UNNAMED is set to 1 (one); otherwise, NAME is set to the <derived column> name for the column and UNNAMED is set to 0 (zero).
  - 3) If the column is a member of the primary key of *T* and KEY\_TYPE was set to 1 (one) or if the column is a member of the preferred candidate key of *T* and KEY\_TYPE was set to 2, then KEY\_MEMBER is set to 1 (one); otherwise, KEY\_MEMBER is set to 0 (zero).
- ii) If the value of LEVEL is 0 (zero) and the item descriptor area describes a <dynamic parameter specification>, then:
- 1) NULLABLE is set to 1 (one).  
NOTE 451 — This indicates that the <dynamic parameter specification> can have the null value.
  - 2) UNNAMED is set to 1 (one) and NAME is set to an implementation-dependent value.
  - 3) KEY\_MEMBER is set to 0 (zero).
- iii) Otherwise:
- 1) NULLABLE is set to 1 (one).
  - 2) Case:
    - A) If the item descriptor area describes a field of a row, then  
Case:
      - I) If the name of the field is implementation-dependent, then NAME is set to the implementation-dependent name of the field and UNNAMED is set to 1 (one).
      - II) Otherwise, NAME is set to the name of the field and UNNAMED is set to 0 (zero).
    - B) Otherwise, UNNAMED is set to 1 (one) and NAME is set to an implementation-defined value.
  - 3) KEY\_MEMBER is set to 0 (zero).
- d) Case:
- i) If TYPE indicates a <character string type>, then:
    - 1) LENGTH is set to the length or maximum length in characters of the character string type.
    - 2) OCTET\_LENGTH is set to the maximum possible length in octets of the character string type.
    - 3) CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA, and CHARACTER\_SET\_NAME are set to the fully qualified name of the character string type's character set.
    - 4) COLLATION\_CATALOG, COLLATION\_SCHEMA and COLLATION\_NAME are set to the fully qualified name of the character string type's declared type collation, if any, and otherwise to the empty string.

If the subject <language clause> specifies C, then the lengths specified in LENGTH and OCTET\_LENGTH do not include the implementation-defined null character that terminates a C character string.

- ii) If TYPE indicates a <binary string type>, then LENGTH is set to the length or maximum length in octets of the binary string and OCTET\_LENGTH is set to the maximum possible length in octets of the binary string.
  - iii) If TYPE indicates an <exact numeric type>, then PRECISION and SCALE are set to the precision and scale of the exact numeric.
  - iv) If TYPE indicates an <approximate numeric type>, then PRECISION is set to the precision of the approximate numeric.
  - v) If TYPE indicates a <datetime type>, then LENGTH is set to the length in positions of the datetime type, DATETIME\_INTERVAL\_CODE is set to a code as specified in Table 26, “Codes associated with datetime data types in Dynamic SQL”, to indicate the specific datetime data type and PRECISION is set to the <time precision> or <timestamp precision>, if either is applicable.
  - vi) If TYPE indicates an <interval type>, then LENGTH is set to the length in positions of the interval type, DATETIME\_INTERVAL\_CODE is set to a code as specified in Table 27, “Codes used for <interval qualifier>s in Dynamic SQL”, to indicate the <interval qualifier> of the interval data type, DATETIME\_INTERVAL\_PRECISION is set to the <interval leading field precision> and PRECISION is set to the <interval fractional seconds precision>, if applicable.
  - vii) If TYPE indicates a user-defined type, then USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME are set to the fully qualified name of the user-defined type, and USER\_DEFINED\_TYPE\_CODE is set to a code as specified in Table 29, “Codes associated with user-defined types in Dynamic SQL”, to indicate the category of the user-defined type.
  - viii) If TYPE indicates a <reference type>, then:
    - 1) USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME are set to the fully qualified name of the referenced type.
    - 2) SCOPE\_CATALOG, SCOPE\_SCHEMA, and SCOPE\_NAME are set to the fully qualified name of the referenceable base table.
    - 3) LENGTH and OCTET\_LENGTH are set to the length in octets of the <reference type>.
  - ix) If TYPE indicates ROW, then DEGREE is set to the degree of the row type.
  - x) 09 If TYPE indicates ARRAY, then CARDINALITY is set to the maximum cardinality of the array type.
- e) If LEVEL is 0 (zero) and the prepared statement is a <call statement>, then:
- i) Let  $SR$  be the subject routine for the <routine invocation> of the <call statement>.
  - ii) Let  $D_x$  be the  $x$ -th <dynamic parameter specification> simply contained in an SQL argument  $A_y$  of the <call statement>.
  - iii) Let  $P_y$  be the  $y$ -th SQL parameter of  $SR$ .

NOTE 452 — A  $P$  whose <parameter mode> is IN can be a <value expression> that contains zero, one, or more <dynamic parameter specification>s. Thus:

- Every  $D_x$  maps to one and only one  $P_y$ .
  - Several  $D_x$  instances can map to the same  $P_y$ .
  - There can be  $P_y$  instances that have no  $D_x$  instances that map to them.
- iv) The PARAMETER\_MODE value in the descriptor for each  $D_x$  is set to the value from Table 28, “Codes used for input/output SQL parameter modes in Dynamic SQL”, that indicates the <parameter mode> of  $P_y$ .
- v) The PARAMETER\_ORDINAL\_POSITION value in the descriptor for each  $D_x$  is set to the ordinal position of  $P_y$ .
- vi) The PARAMETER\_SPECIFIC\_CATALOG, PARAMETER\_SPECIFIC\_SCHEMA, and PARAMETER\_SPECIFIC\_NAME values in the descriptor for each  $D_x$  are set to the values that identify the catalog, schema, and specific name of  $SR$ .

## Conformance Rules

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <describe input statement>.
- 2) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <describe output statement>.



## 20.10 <input using clause>

*This Subclause is modified by Subclause 15.2, “<input using clause>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 17.5, “<input using clause>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 17.2, “<input using clause>”, in ISO/IEC 9075-14.*

### Function

Supply input values for an <SQL dynamic statement>.

### Format

```
<input using clause> ::=
 <using arguments>
 | <using input descriptor>

<using arguments> ::=
 USING <using argument> [{ <comma> <using argument> }...]

<using argument> ::=
 <general value specification>

<using input descriptor> ::=
 <using descriptor>
```

### Syntax Rules

- 1) 04 The <general value specification> immediately contained in <using argument> shall be either a <host parameter specification>, an <SQL parameter reference>, or an <embedded variable specification>.

### Access Rules

*None.*

### General Rules

- 1) 09 If <descriptor name> does not identify an SQL descriptor area, then an exception condition is raised: *invalid SQL descriptor name*.
- 2) When an <input using clause> is used in a <dynamic open statement> or as the <parameter using clause> in an <execute statement>, the <input using clause> describes the input <dynamic parameter specification> values for the <dynamic open statement> or the <execute statement>, respectively. Let *PS* be the prepared <dynamic select statement> referenced by the <dynamic open statement> or the prepared statement referenced by the <execute statement>, respectively.
- 3) Let *D* be the number of input <dynamic parameter specification>s in *PS*.
- 4) If <using arguments> is specified and the number of <using argument>s is not *D*, then an exception condition is raised: *dynamic SQL error — using clause does not match dynamic parameter specifications*.

5) If <using input descriptor> is specified, then:

- a) Let  $N$  be the value of COUNT.
- b) If  $N$  is greater than the value of <occurrences> specified when the SQL descriptor area identified by <descriptor name> was allocated or is less than zero, then an exception condition is raised: *dynamic SQL error — invalid descriptor count*.
- c) If the first  $N$  item descriptor areas are not valid as specified in Subclause 20.1, “Description of SQL descriptor areas”, then an exception condition is raised: *dynamic SQL error — using clause does not match dynamic parameter specifications*.
- d) In the first  $N$  item descriptor areas:
  - i) If the number of item descriptor areas in which the value of LEVEL is 0 (zero) is not  $D$ , then an exception condition is raised: *dynamic SQL error — using clause does not match dynamic parameter specifications*.
  - ii) If the value of INDICATOR is not negative, TYPE does not indicate ROW, and the item descriptor area is not subordinate to an item descriptor area whose INDICATOR value is negative or whose TYPE field indicates ARRAY, ARRAY LOCATOR, MULTISSET, or MULTISSET LOCATOR, and if the value of DATA is not a valid value of the data type represented by the item descriptor area, then an exception condition is raised: *dynamic SQL error — using clause does not match dynamic parameter specifications*.

6) For 1 (one)  $\leq i \leq D$ :

- a) Let  $TDT$  be the effective declared type of the  $i$ -th input <dynamic parameter specification>, defined to be the type represented by the item descriptor area and its subordinate descriptor areas that would be set by a <describe input statement> to reflect the description of the  $i$ -th input <dynamic parameter specification> of  $PS$ .

NOTE 453 — See the General Rules of Subclause 20.9, “<describe statement>”.

NOTE 454 — “Represented by”, as applied to the relationship between a data type and an item descriptor area, is defined in the Syntax Rules of Subclause 20.1, “Description of SQL descriptor areas”.

b) Case:

- i) If <using input descriptor> is specified, then:

- 1) Let  $IDA$  be the  $i$ -th item descriptor area whose LEVEL value is 0 (zero).
- 2) Let  $SDT$  be the effective declared type represented by  $IDA$ .

NOTE 455 — “Represented by”, as applied to the relationship between a data type and an item descriptor area, is defined in the Syntax Rules of Subclause 20.1, “Description of SQL descriptor areas”.

- 3) Let  $SV$  be the *associated value* of  $IDA$ .

Case:

A) If the value of INDICATOR is negative, then  $SV$  is the null value.

B) Otherwise,

Case:

- I) If TYPE indicates ROW, then *SV* is a row whose type is *SDT* and whose field values are the associated values of the immediately subordinate descriptor areas of *IDA*.
  - II) Otherwise, *SV* is the value of DATA with data type *SDT*.
- ii) If <using arguments> is specified, then let *SDT* and *SV* be the declared type and value, respectively, of the *i*-th <using argument>.
- c) Case:
- i) If *SDT* is a locator type, then  
Case:
    - 1) If *SV* is not the null value, then let the value of the *i*-th dynamic parameter be the value of *SV*.
    - 2) Otherwise, let the value of the *i*-th dynamic parameter be the null value.
  - ii) 14 If *SDT* and *TDT* are predefined data types, then  
Case:
    - 1) If the <cast specification>  

`CAST ( SV AS TDT )`

does not conform to the Syntax Rules of Subclause 6.12, “<cast specification>”, and there is an implementation-defined conversion from type *STD* to type *TDT*, then that implementation-defined conversion is effectively performed, converting *SV* to type *TDT*, and the result is the value *TV* of the *i*-th input dynamic parameter.
    - 2) Otherwise:
      - A) If the <cast specification>  

`CAST ( SV AS TDT )`

does not conform to the Syntax Rules of Subclause 6.12, “<cast specification>”, then an exception condition is raised: *dynamic SQL error — restricted data type attribute violation*.
      - B) The <cast specification>  

`CAST ( SV AS TDT )`

is effectively performed and is the value of the *i*-th input dynamic parameter.
  - iii) If *SDT* is a predefined data type and *TDT* is a user-defined type, then:
    - 1) Let *DT* be the data type identified by *TDT*.
    - 2) If the current SQL-session has a group name corresponding to the user-defined type name of *DT*, then let *GN* be that group name; Otherwise, let *GN* be the default transform group name associated with the current SQL-session.

- 3) The Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, are applied with *DT* and *GN* as *TYPE* and *GROUP*, respectively.

Case:

- A) If there is an applicable to-sql function, then let *TSF* be that to-sql function. If *TSF* is an SQL-invoked method, then let *TSFPT* be the declared type of the second SQL parameter of *TSF*; otherwise, let *TSFPT* be the declared type of the first SQL parameter of *TSF*.

Case:

- I) If *TSFPT* is compatible with *SDT*, then

Case:

- 1) If *TSF* is an SQL-invoked method, then *TSF* is effectively invoked with the value returned by the function invocation:

*DT* ( )

as the first parameter and *SV* as the second parameter. The <return value> is the value of the *i*-th input dynamic parameter.

- 2) Otherwise, *TSF* is effectively invoked with *SV* as the first parameter. The <return value> is the value of the *i*-th input dynamic parameter.

- II) Otherwise, an exception condition is raised: *dynamic SQL error — restricted data type attribute violation*.

- B) Otherwise, an exception condition is raised: *dynamic SQL error — data type transform function violation*.

## Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <input using clause>.

## 20.11 <output using clause>

*This Subclause is modified by Subclause 15.3, “<output using clause>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 17.6, “<output using clause>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 17.3, “<output using clause>”, in ISO/IEC 9075-14.*

### Function

Supply output variables for an <SQL dynamic statement>.

### Format

```
<output using clause> ::=
 <into arguments>
 | <into descriptor>

<into arguments> ::=
 INTO <into argument> [{ <comma> <into argument> }...]

<into argument> ::=
 <target specification>

<into descriptor> ::=
 INTO [SQL] DESCRIPTOR <descriptor name>
```

### Syntax Rules

- 1) 04 The <target specification> immediately contained in <into argument> shall be either a <host parameter specification>, an <SQL parameter reference>, or an <embedded variable specification>.

### Access Rules

*None.*

### General Rules

- 1) 09 When an <output using clause> is used in a <dynamic fetch statement> or as the <result using clause> of an <execute statement>, let *PS* be the prepared <dynamic select statement> referenced by the <dynamic fetch statement> or the prepared <dynamic single row select statement> referenced by the <execute statement>, respectively.
- 2) Case:
  - a) If *PS* is a <dynamic select statement> or a <dynamic single row select statement>, then let *D* be the degree of the table specified by *PS*.
  - b) Otherwise, let *D* be the number of output <dynamic parameter specification>s contained in *PS*.
- 3) If <into arguments> is specified and the number of <into argument>s is not *D*, then an exception condition is raised: *dynamic SQL error — using clause does not match target specifications*.

- 4) If <into descriptor> is specified, then:
- If <descriptor name> does not identify an SQL descriptor area, then an exception condition is raised: *invalid SQL descriptor name*.
  - Let  $N$  be the value of COUNT.
  - If  $N$  is greater than the value of <occurrences> specified when the SQL descriptor area identified by <descriptor name> was allocated or less than zero, then an exception condition is raised: *dynamic SQL error — invalid descriptor count*.
  - If the first  $N$  item descriptor areas are not valid as specified in Subclause 20.1, “Description of SQL descriptor areas”, then an exception condition is raised: *dynamic SQL error — using clause does not match target specifications*.
  - In the first  $N$  item descriptor areas, if the number of item descriptor areas in which the value of LEVEL is 0 (zero) is not  $D$ , then an exception condition is raised: *dynamic SQL error — using clause does not match target specifications*.
- 5) For  $1 \text{ (one)} \leq i \leq D$ :
- Let  $SDT$  be the effective declared type of the  $i$ -th <select list> column or output <dynamic parameter specification>, defined to be the type represented by the item descriptor area and its subordinate descriptor areas that would be set by  
Case:  
    - If  $PS$  is a <dynamic select statement> or a <dynamic single row select statement>, then a <describe output statement> to reflect the description of the  $i$ -th <select list> column; let  $SV$  be the value of that <select list> column, with data type  $SDT$ .
    - Otherwise, a <describe output statement> to reflect the description of the  $i$ -th output <dynamic parameter specification>; let  $SV$  be the value of that <dynamic parameter specification>, with data type  $SDT$ .

NOTE 456 — “Represented by”, as applied to the relationship between a data type and an item descriptor area, is defined in the Syntax Rules of Subclause 20.1, “Description of SQL descriptor areas”.
  - Case:  
    - If <into descriptor> is specified, then let  $TDT$  be the declared type of the  $i$ -th <target specification> as represented by the  $i$ -th item descriptor area  $IDA$  whose LEVEL value is 0 (zero).  

NOTE 457 — “Represented by”, as applied to the relationship between a data type and an item descriptor area, is defined in the Syntax Rules of Subclause 20.1, “Description of SQL descriptor areas”.
    - If <into arguments> is specified, then let  $TDT$  be the data type of the  $i$ -th <into argument>.
  - If the <output using clause> is directly contained in a <dynamic fetch statement>, then let  $CR$  be the dynamic cursor identified by the <dynamic fetch statement>, and let  $LTDT$  be the most specific type of the  $i$ -th <target specification> or <into argument> of the most recently executed <dynamic fetch statement> prior to the current execution, if any, for the cursor  $CR$ . It is implementation-defined whether or not an exception condition is raised: *dynamic SQL error — restricted data type attribute violation* if any of the following are true:
    - $LTDT$  and  $TDT$  both identify a binary large object type and only one of  $LTDT$  and  $TDT$  is a binary large object locator.

- ii) *LTDT* and *TDT* both identify a character large object type and only one of *LTDT* and *TDT* is a character large object locator.
  - iii) *LTDT* and *TDT* both identify an array type and only one of *LTDT* and *TDT* is an array locator.
  - iv) *LTDT* and *TDT* both identify a multiset type and only one of *LTDT* and *TDT* is a multiset locator.
  - v) *LTDT* and *TDT* both identify a user-defined type and only one of *LTDT* and *TDT* is a user-defined type locator.
- d) Case:
- i) If *TDT* is a locator type, then  
Case:
    - 1) If *SV* is not the null value, then a locator *L* that uniquely identifies *SV* is generated and is the value *TV* of the *i*-th <target specification>.
    - 2) Otherwise, the value *TV* of the *i*-th <target specification> is the null value.
  - ii) 14 If *STD* and *TDT* are predefined data types, then  
Case:
    - 1) If the <cast specification>  

*CAST* ( *SV* AS *TDT* )

does not conform to the Syntax Rules of Subclause 6.12, “<cast specification>”, and there is an implementation-defined conversion of type *STD* to type *TDT*, then that implementation-defined conversion is effectively performed, converting *SV* to type *TDT*, and the result is the value *TV* of the *i*-th <target specification>.
    - 2) Otherwise:
      - A) If the <cast specification>  

*CAST* ( *SV* AS *TDT* )

does not conform to the Syntax Rules of Subclause 6.12, “<cast specification>”, then an exception condition is raised: *dynamic SQL error — restricted data type attribute violation*.
      - B) The <cast specification>  

*CAST* ( *SV* AS *TDT* )

is effectively performed, and is the value *TV* of the *i*-th <target specification>.
  - iii) If *SDT* is a user-defined type and *TDT* is a predefined data type, then:
    - 1) Let *DT* be the data type identified by *SDT*.
    - 2) If the current SQL-session has a group name corresponding to the user-defined type name of *DT*, then let *GN* be that group name; otherwise, let *GN* be the default transform group name associated with the current SQL-session.

- 3) Apply the Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, with *DT* and *GN* as *TYPE* and *GROUP*, respectively.

Case:

- A) If there is an applicable from-sql function, then let *FSF* be that from-sql function and let *FSFRT* be the <returns data type> of *FSF*.

Case:

- I) If *FSFRT* is compatible with *TDT*, then the from-sql function *FSF* is effectively invoked with *SV* as its input SQL parameter and the <return value> is the value *TV* of the *i*-th <target specification>.
- II) Otherwise, an exception condition is raised: *dynamic SQL error — restricted data type attribute violation*.
- B) Otherwise, an exception condition is raised: *dynamic SQL error — data type transform function violation*.

e) Case:

- i) If <into descriptor> is specified, then *IDA* is set to reflect the value of *TV* as follows.

Case:

- 1) If *TYPE* indicates ROW, then

Case:

- A) If *TV* is the null value, then the value of INDICATOR in *IDA* and in all subordinate descriptor areas of *IDA* that are not subordinate to an item descriptor area whose *TYPE* indicates ARRAY, ARRAY LOCATOR, MULTISSET, or MULTISSET LOCATOR is set to -1.
- B) Otherwise, the *i*-th subordinate descriptor area of *IDA* is set to reflect the value of the *i*-th field of *TV* by applying this subrule (beginning with the outermost 'Case') to the *i*-th subordinate descriptor area of *IDA* as *IDA*, the value of the *i*-th field of *TV* as *TV*, the value of the *i*-th field of *SV* as *SV*, and the data type of the *i*-th field of *SV* as *SDT*.

- 2) Otherwise,

Case:

- A) If *TV* is the null value, then the value of INDICATOR is set to -1.

- B) If *TV* is not the null value, then:

- I) The value of INDICATOR is set to 0 (zero).

- II) Case:

- 1) If *TYPE* indicates a locator type, then a locator *L* that uniquely identifies *TV* is generated and the value of DATA is set to an implementation-dependent four-octet integer value that represents *L*.

- 2) Otherwise, the value of DATA is set to *TV*.

- III) Case:



- 1) If TYPE indicates CHARACTER VARYING, CHARACTER LARGE OBJECT, BINARY VARYING, or BINARY LARGE OBJECT, then RETURNED\_LENGTH is set to the length in characters or octets, respectively, of *TV*, and RETURNED\_OCTET\_LENGTH is set to the length in octets of *TV*.
  - 2) If *SDT* is CHARACTER VARYING, CHARACTER LARGE OBJECT, BINARY VARYING, BINARY LARGE OBJECT, then RETURNED\_LENGTH is set to the length in characters or octets, respectively, of *SV*, and RETURNED\_OCTET\_LENGTH is set to the length in octets of *SV*.
  - 3) If TYPE indicates ARRAY, ARRAY LOCATOR, MULTISSET, or MULTISET LOCATOR, then RETURNED\_CARDINALITY is set to the cardinality of *TV*.
- ii) If <into arguments> is specified, then the Rules in Subclause 9.1, “Retrieval assignment”, are applied to *TV* and the *i*-th <into argument> as *VALUE* and *TARGET*, respectively.

NOTE 458 — All other values of the SQL descriptor area are unchanged.

## Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <output using clause>.

## 20.12 <execute statement>

This Subclause is modified by *Subclause 17.7*, “<execute statement>”, in ISO/IEC 9075-9.

### Function

Associate input SQL parameters and output targets with a prepared statement and execute the statement.

### Format

```
<execute statement> ::=
 EXECUTE <SQL statement name> [<result using clause>] [<parameter using clause>]

<result using clause> ::=
 <output using clause>

<parameter using clause> ::=
 <input using clause>
```

### Syntax Rules

- 1) If <SQL statement name> is a <statement name>, then

Case:

- a) If the <execute statement> is contained in an <SQL-invoked routine>, then the innermost containing <SQL-invoked routine> shall contain a <prepare statement> whose <statement name> is equivalent to the <statement name> of the <execute statement>.
- b) Otherwise, the containing <SQL-client module definition> shall contain a <prepare statement> whose <statement name> is equivalent to the <statement name> of the <execute statement>.

### Access Rules

*None.*

### General Rules

- 1) 09 If the <SQL statement name> does not identify a prepared statement *P*, then an exception condition is raised: *invalid SQL statement name*.
- 2) Let *PS* be the statement previously prepared using <SQL statement name>.
- 3) If *PS* is a <dynamic select statement>, then:
  - a) If *PS* does not conform to the Format and Syntax Rules of a <dynamic single row select statement>, then an exception condition is raised: *dynamic SQL error — cursor specification cannot be executed*.
  - b) Otherwise, *PS* is treated as a <dynamic single row select statement>.

- 4) If *PS* contains the <table name> of a created or declared local temporary table and if the <execute statement> is not in the same <SQL-client module definition> as the <prepare statement> that prepared the prepared statement, then an exception condition is raised: *syntax error or access rule violation*.
- 5) If *PS* contains input <dynamic parameter specification>s and a <parameter using clause> is not specified, then an exception condition is raised: *dynamic SQL error — using clause required for dynamic parameters*.
- 6) If *PS* is a <dynamic single row select statement> or it contains output <dynamic parameter specification>s and a <result using clause> is not specified, then an exception condition is raised: *dynamic SQL error — using clause required for result fields*.
- 7) If a <parameter using clause> is specified, then the General Rules specified in Subclause 20.10, “<input using clause>”, for a <parameter using clause> in an <execute statement> are applied.
- 8) A copy of the top cell is pushed onto the authorization stack. If *PS* has an owner, then the top cell of the authorization stack is set to contain only the authorization identifier of the owner of *PS*.
- 9) The General Rules of Subclause 13.5, “<SQL procedure statement>”, are evaluated with *PS* as the executing statement.
- 10) If a <result using clause> is specified, then the General Rules specified in Subclause 20.11, “<output using clause>”, for a <result using clause> in an <execute statement> are applied.
- 11) Upon completion of execution, the top cell in the authorization stack is removed.

## Conformance Rules

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <result using clause>.
- 2) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <execute statement>.

## 20.13 <execute immediate statement>

### Function

Dynamically prepare and execute a preparable statement.

### Format

```
<execute immediate statement> ::=
 EXECUTE IMMEDIATE <SQL statement variable>
```

### Syntax Rules

- 1) The declared type of <SQL statement variable> shall be character string.

### Access Rules

*None.*

### General Rules

- 1) Let *P* be the contents of the <SQL statement variable>.
- 2) If one or more of the following are true, then an exception condition is raised: *syntax error or access rule violation*.
  - a) *P* is a <dynamic select statement> or a <dynamic single row select statement>.
  - b) *P* contains a <dynamic parameter specification>.
- 3) Let *SV* be <SQL statement variable>. <execute immediate statement> is equivalent to the following:

```
PREPARE IMMEDIATE_STMT FROM SV ;
EXECUTE IMMEDIATE_STMT ;
DEALLOCATE PREPARE IMMEDIATE_STMT ;
```

where *IMMEDIATE\_STMT* is an implementation-dependent <statement name> that does not identify any existing prepared statement.

NOTE 459 — Exception condition or completion condition information resulting from the PREPARE or EXECUTE is reflected in the diagnostics area.

### Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <execute immediate statement>.

## 20.14 <dynamic declare cursor>

*This Subclause is modified by Subclause 17.8, “<dynamic declare cursor>”, in ISO/IEC 9075-9.*

### Function

Declare a declared dynamic cursor to be associated with a <statement name>, which may in turn be associated with a <cursor specification>.

### Format

```
<dynamic declare cursor> ::=
 DECLARE <cursor name>
 <cursor properties>
 FOR <statement name>
```

### Syntax Rules

- 1) The <cursor name> shall not be identical to the <cursor name> specified in any other <declare cursor> or <dynamic declare cursor> in the same <SQL-client module definition> *M*. The scope of the <cursor name> is *M* with the exception of any <SQL schema statement> contained in *M*.
- 2) Let *SN* be the <statement name> simply contained in the <dynamic declare cursor>. The containing <SQL-client module definition> shall contain, without an intervening <SQL schema statement>, a <prepare statement> whose <statement name> is equivalent to *SN*.

### Access Rules

*None.*

### General Rules

- 1) 09 A cursor declaration descriptor *CDD* is created. *CDD* includes indications that:
  - a) The kind of cursor is a declared dynamic cursor.
  - b) The provenance of the cursor is an indication of the SQL-client module whose <SQL-client module definition> contains the <dynamic declare cursor>.
  - c) The name of the cursor is the <cursor name>.
  - d) The cursor's origin is *SN*.
  - e) The cursor's declared properties are as determined by the <cursor properties>.

### Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic declare cursor>.

## 20.15 <allocate cursor statement>

This Subclause is modified by *Subclause 17.9*, “<allocate cursor statement>”, in ISO/IEC 9075-9.

### Function

Define a cursor based on a prepared statement for a <cursor specification> or assign a cursor to the result set sequence returned from an SQL-invoked procedure.

### Format

```
<allocate cursor statement> ::=
 ALLOCATE <extended cursor name> <cursor intent>

<cursor intent> ::=
 <statement cursor>
 | <result set cursor>

<statement cursor> ::=
 <cursor properties>
 FOR <extended statement name>

<result set cursor> ::=
 [CURSOR] FOR PROCEDURE <specific routine designator>
```

### Syntax Rules

- 1) If <result set cursor> is specified, then the SQL-invoked routine identified by <specific routine designator> shall be an SQL-invoked procedure.

### Access Rules

*None.*

### General Rules

- 1) 09 Let *S* be the <simple value specification> immediately contained in <extended cursor name>. Let *V* be the character string that is the result of

```
TRIM (BOTH ' ' FROM S)
```

Case:

- a) If *V* does not conform to the Format and Syntax Rules of an <identifier>, then an exception condition is raised: *invalid cursor name*.
  - b) Otherwise, let *ECN* be the <extended statement name>. The value of *ECN* is *V*.
- 2) If *ECN* identifies a cursor, then an exception condition is raised: *invalid cursor name*.
  - 3) If <statement cursor> is specified, then:

- a) If <extended statement name> does not identify a prepared statement, then an exception condition is raised: *invalid SQL statement name*.
  - b) If the prepared statement *P* identified by <extended statement name> is not a <cursor specification>, then an exception condition is raised: *dynamic SQL error — prepared statement not a cursor specification*.
  - c) *P* is re-prepared as follows: if *P* does not conform to the Syntax Rules of Subclause 14.3, “<cursor specification>”, then an exception condition is raised: *syntax error or access rule violation*.
  - d) A cursor declaration descriptor *CDD* is created. *CDD* includes:
    - i) The kind of cursor is an extended statement cursor.
    - ii) The provenance of the cursor is  
Case:
      - 1) If the <scope option> contained in <extended cursor name> is GLOBAL, then the current SQL-session identifier.
      - 2) Otherwise, an indication of the SQL-client module whose <SQL-client module definition> contains the <allocate cursor statement>.
    - iii) The name of the cursor is the extended name *V* and the explicit or implicit <scope option> of the <extended cursor name>.
    - iv) The cursor origin is *P*.
    - v) The cursor's declared properties are as determined by the <cursor properties>.
  - e) A cursor instance descriptor *CID* is created. *CID* includes:
    - i) The cursor declaration descriptor is *CDD*.
    - ii) The SQL-session identifier is the current SQL-session identifier.
    - iii) The cursor's state is closed.
  - f) An association is made between the value of the <extended cursor name> and the prepared statement in the scope of the <extended cursor name>. The association is preserved until the prepared statement is destroyed, at which time cursor declaration descriptor and the cursor instance descriptor of the cursor identified by <extended cursor name> are also destroyed.
- 4) If <result set cursor> is specified, then:
- a) Let *SIP* be the SQL-invoked procedure identified by <specific routine designator>. Let *INV* be the active SQL-invoked routine of the current routine execution context.
  - b) If the SQL-session context of the current SQL-session does not include a result set sequence *RSS* brought into existence by an invocation of *SIP* by *INV*, then an exception condition is raised: *invalid SQL-invoked procedure reference*.
  - c) If *RSS* is empty, then an exception condition is raised: *no data — no additional result sets returned*.
  - d) A cursor declaration descriptor *CDD* is created. *CDD* includes indications that:
    - i) The kind of cursor is a received cursor.

- ii) The provenance of the cursor is  
Case:
  - 1) If the <scope option> contained in <extended cursor name> is GLOBAL, then the current SQL-session identifier.
  - 2) Otherwise, an indication of the SQL-client module whose <SQL-client module definition> contains the <allocate cursor statement>.
- iii) The name of the cursor is the extended name *V* and the explicit or implicit <scope option> of the <extended cursor name>.
- iv) The cursor origin is the <specific routine designator>.
- v) The cursor's declared properties are as follows:
  - 1) The cursor's declared sensitivity is ASENSITIVE.
  - 2) The cursor's declared scrollability is NO SCROLL.
  - 3) The cursor's declared holdability is WITHOUT HOLD.
  - 4) The cursor's declared returnability is WITHOUT RETURN.
- e) A cursor instance descriptor *CID* is created. *CID* includes:
  - i) The cursor declaration descriptor is *CDD*.
  - ii) The current SQL-session identifier.
  - iii) The cursor's state is open.
- f) The General Rules of Subclause 15.2, “Effect of receiving a result set”, are applied, with *CID* as *CURSOR* and *RSS* as *RESULT SET SEQUENCE*.

## Conformance Rules

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain an <allocate cursor statement>.
- 2) Without Feature T471, “Result sets return value”, conforming SQL language shall not contain a <result set cursor>.



## 20.16 <dynamic open statement>

This Subclause is modified by [Subclause 17.10](#), “<dynamic open statement>”, in ISO/IEC 9075-9.

### Function

Associate input dynamic parameters with a <cursor specification> and open the dynamic cursor.

### Format

```
<dynamic open statement> ::=
 OPEN <dynamic cursor name> [<input using clause>]
```

### Syntax Rules

- 1) If <dynamic cursor name> *DCN* is a <cursor name> *CN*, then *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*. *CN* shall identify a declared dynamic cursor. Let *CDD* be the cursor declaration descriptor identified by *CN*.

### Access Rules

- 1) The Access Rules for the <query expression> simply contained in the prepared statement associated with the <dynamic cursor name> are applied.

### General Rules

- 1) Case:
  - a) 09 If *DCN* is a <cursor name> *CN*, then  
Case:
    - i) If the <statement name> *SN* contained in *CDD* does not identify a prepared statement, then an exception condition is raised: *invalid SQL statement name*.
    - ii) If *SN* does not identify a prepared statement that is a <cursor specification>, then an exception condition is raised: *dynamic SQL error — prepared statement not a cursor specification*.
  - b) Otherwise, if *DCN* does not identify an extended statement cursor, then an exception condition is raised: *invalid cursor name*.
- 2) Let *CR* be the cursor identified by *DCN*.
- 3) If the prepared statement *P* associated with the <dynamic cursor name> contains <dynamic parameter specification>s and an <input using clause> is not specified, then an exception condition is raised: *dynamic SQL error — using clause required for dynamic parameters*.
- 4) If the <dynamic cursor name> is a <cursor name>, then *P* is re-prepared as follows: if *P* does not conform to the Syntax Rules of [Subclause 14.3](#), “<cursor specification>”, then an exception condition is raised: *syntax error or access rule violation*.

**20.16 <dynamic open statement>**

- 5) *CR* is updatable if and only if the <cursor specification> included in the result set descriptor of *CR* is updatable.
- 6) If an <input using clause> is specified, then the General Rules specified in Subclause 20.10, “<input using clause>”, for <dynamic open statement> are applied.
- 7) The General Rules of Subclause 15.1, “Effect of opening a cursor”, are applied with *CR* as *CURSOR*.

## **Conformance Rules**

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic open statement>.

## 20.17 <dynamic fetch statement>

This Subclause is modified by *Subclause 17.11*, “<dynamic fetch statement>”, in ISO/IEC 9075-9.

### Function

Fetch a row for a dynamic cursor.

### Format

```
<dynamic fetch statement> ::=
 FETCH [[<fetch orientation>] FROM] <dynamic cursor name> <output using clause>
```

### Syntax Rules

- 1) If <fetch orientation> is omitted, then NEXT is implicit.
- 2) If <dynamic cursor name> *DCN* is a <cursor name> *CN*, then *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*. *CN* shall identify a declared dynamic cursor.

### Access Rules

*None.*

### General Rules

- 1) 09 If *DCN* does not identify a cursor, then an exception condition is raised: *invalid cursor name*.
- 2) Let *CR* be the cursor identified by *DCN*.
- 3) The General Rules of [Subclause 15.3](#), “Determination of the current row of a cursor”, are applied with *CR* as *CURSOR* and <fetch orientation> as *FETCH ORIENTATION*.
- 4) If a completion condition *no data* is raised, then no further General Rules of this Subclause are applied.
- 5) The General Rules specified in [Subclause 20.11](#), “<output using clause>”, for an <output using clause> in a <dynamic fetch statement> are applied.
- 6) If an exception condition is raised during the assignment of a value to a target, then the values of all targets are implementation-dependent and *CR* remains positioned on the current row.

NOTE 460 — It is implementation-dependent whether *CR* remains positioned on the current row when an exception condition is raised during the derivation of any <derived column>.

### Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic fetch statement>.

## 20.18 <dynamic single row select statement>

### Function

Retrieve values from a dynamically-specified row of a table.

### Format

```
<dynamic single row select statement> ::=
 <query specification>
```

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let  $Q$  be the result of the <query specification>.
- 2) Case:
  - a) If the cardinality of  $Q$  is greater than 1 (one), then an exception condition is raised: *cardinality violation*.
  - b) If  $Q$  is empty, then a completion condition is raised: *no data*.

### Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic single row select statement>.

## 20.19 <dynamic close statement>

This Subclause is modified by [Subclause 17.12](#), “<dynamic close statement>”, in ISO/IEC 9075-9.

### Function

Close a dynamic cursor.

### Format

```
<dynamic close statement> ::=
 CLOSE <dynamic cursor name>
```


### Syntax Rules

- 1) If <dynamic cursor name> *DCN* is a <cursor name> *CN*, then *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*. *CN* shall identify a declared dynamic cursor.

### Access Rules

*None.*

### General Rules

- 1) If *DCN* does not identify a cursor, then an exception condition is raised: *invalid cursor name*.
- 2) Let *CR* be the cursor identified by *DCN*.
- 3)  The General Rules of [Subclause 15.4](#), “Effect of closing a cursor”, are applied with *CR* as *CURSOR* and DESTROY as *DISPOSITION*.

### Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic close statement>.

## 20.20 <dynamic delete statement: positioned>

### Function

Delete a row of a table.

### Format

```
<dynamic delete statement: positioned> ::=
 DELETE FROM <target table> WHERE CURRENT OF <dynamic cursor name>
```

### Syntax Rules

- 1) If <dynamic cursor name> *DCN* is a <cursor name> *CN*, then *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*. *CN* shall identify a declared dynamic cursor.
- 2) Let *TN* be the <table name> contained in <target table> *TT*.

### Access Rules

- 1) Case:
  - a) If <dynamic delete statement: positioned> *DDSP* is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges for the owner of that schema shall include DELETE for *TN*.
  - b) Otherwise, the current privileges shall include DELETE for *TN*.

### General Rules

- 1) If *DCN* does not identify a dynamic cursor, then an exception condition is raised: *invalid cursor name*.
- 2) Let *CR* be the cursor identified by *DCN*.
- 3) If *CR* is not an updatable cursor, then an exception condition is raised: *invalid cursor name*.
- 4) Let *T* be the simply underlying table of *CR*. *T* is the subject table of the <dynamic delete statement: positioned>. Let *LUT* be the leaf underlying table of *T* such that *T* is one-to-one with *LUT*. Let *LUTN* be a <table name> that identifies *LUT*.
- 5) If *TN* does not identify *LUTN*, or if ONLY is specified and the <table reference> in *T* that references *LUT* does not specify ONLY, or if ONLY is not specified and the <table reference> in *T* that references *LUT* does specify ONLY, then an exception condition is raised: *target table disagrees with cursor specification*.
- 6) The General Rules of Subclause 15.5, “Effect of a positioned delete”, are applied with *CR* as *CURSOR*, *DDSP* as *STATEMENT*, and *TT* as *TARGET*.

## Conformance Rules

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic delete statement: positioned>.

## 20.21 <dynamic update statement: positioned>

### Function

Update a row of a table.

### Format

```
<dynamic update statement: positioned> ::=
 UPDATE <target table> SET <set clause list>
 WHERE CURRENT OF <dynamic cursor name>
```

### Syntax Rules

- 1) If <dynamic cursor name> *DCN* is a <cursor name> *CN*, then *CN* shall be contained within the scope of a <cursor name> that is equivalent to *CN*. *CN* shall identify a declared dynamic cursor.
- 2) Let *TN* be the <table name> contained in <target table> *TT*.
- 3) The scope of *TN* is the entire <dynamic update statement: positioned> *DUSP*.

### Access Rules

- 1) Case:
  - a) If <dynamic update statement: positioned> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema. The applicable privileges for *A* shall include UPDATE for each <object column>.
  - b) Otherwise, the current privileges shall include UPDATE for each <object column>.

### General Rules

- 1) If *DCN* does not identify a dynamic cursor, then an exception condition is raised: *invalid cursor name*.
- 2) Let *CR* be the cursor identified by *DCN*.
- 3) If *CR* is not an updatable cursor, then an exception condition is raised: *invalid cursor name*.
- 4) Let *T* be the simply underlying table of *CR*. *T* is the subject table of the <dynamic update statement: positioned>. Let *LUT* be the leaf underlying table of *T* such that *T* is one-to-one with *LUT*. Let *LUTN* be a <table name> that identifies *LUT*.
- 5) If *TN* does not identify *LUTN*, or if ONLY is specified and the <table reference> in *T* that references *LUT* does not specify ONLY, or if ONLY is not specified and the <table reference> in *T* that references *LUT* does specify ONLY, then an exception condition is raised: *target table disagrees with cursor specification*.
- 6) If any object column is directly or indirectly referenced in the <order by clause> simply contained in the <cursor specification> for *CR*, then an exception condition is raised: *attempt to assign to ordering column*.



**20.21 <dynamic update statement: positioned>**

- 7) If any object column identifies a column that is not identified by a <column name> contained in the explicit or implicit <column name list> of the explicit or implicit <updatability clause> of the <cursor specification> included in the result set descriptor of *CR*, then an exception condition is raised: *attempt to assign to non-updatable column*.
- 8) Let *SCL* be the <set clause list>.
- 9) The General Rules of Subclause 15.6, “Effect of a positioned update”, are applied with *CR* as *CURSOR*, *SCL* as *SET CLAUSE LIST*, *DUSP* as *STATEMENT*, and *TT* as *TARGET*.

**Conformance Rules**

- 1) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic update statement: positioned>.

## 20.22 <preparable dynamic delete statement: positioned>

### Function

Delete a row of a table through a dynamic cursor.

### Format

```
<preparable dynamic delete statement: positioned> ::=
 DELETE [FROM <target table>]
 WHERE CURRENT OF <preparable dynamic cursor name>
```

### Syntax Rules

- 1) If the <preparable dynamic cursor name> *PDCN* is not ambiguous or invalid, then:
  - a) Let *CR* be the cursor identified by *PDCN*. *CR* shall be an updatable cursor.
  - b) If <target table> is not specified, then let *QE* be the <query expression> simply contained in the <cursor specification> of the result set descriptor of *CR*. Let *LUT* be the leaf underlying table of *QE* such that *QE* is one-to-one with respect to *LUT*. Let *TN* be the name of *LUT*.

Case:

- i) If the <table reference> that references *LUT* specifies **ONLY**, then the <target table>
 

*ONLY* ( *TN* )

is implicit.
- ii) Otherwise, the <target table>
 

*TN*

is implicit.
- c) Let *TT* be the explicit or implicit <target table>.
- d) Let *TN* be the <table name> contained in *TT*.
- e) If *TT* immediately contains **ONLY** and the table identified by *TN* is not a typed table, then *TT* is equivalent to *TN*.
- f) Let *T* be the simply underlying table of *CR*. *T* is the *subject table* of the <preparable dynamic delete statement: positioned> *PDDSP*. Let *LUT* be the leaf underlying table of *T* such that *T* is one-to-one with respect to *LUT*.
- g) The subject table of a <preparable dynamic delete statement: positioned> shall not identify an old transition table or a new transition table.
- h) *TN* shall identify *LUT*.
- i) <target table> shall specify **ONLY** if and only if the <table reference> contained in *T* that references *LUT* specifies **ONLY**.

**20.22 <preparable dynamic delete statement: positioned>**

- j) The schema identified by the explicit or implicit qualifier of *TN* shall include the descriptor of *LUT*.

**Access Rules**

- 1) If *PDCN* is not ambiguous or invalid, then

Case:

- a) If <preparable dynamic delete statement: positioned> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then the applicable privileges for the owner of that schema shall include DELETE for *TN*.
- b) Otherwise, the current privileges shall include DELETE for *TN*.

**General Rules**

- 1) The General Rules of Subclause 15.5, “Effect of a positioned delete”, are applied with *CR* as *CURSOR*, *PDDSP* as *STATEMENT*, and *TT* as *TARGET*.

NOTE 461 — If the General Rules are reached, then *PDCN* cannot be ambiguous or invalid.

**Conformance Rules**

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <preparable dynamic delete statement: positioned>.

## 20.23 <preparable dynamic cursor name>

This Subclause is modified by *Subclause 8.1*, “<preparable dynamic cursor name>”, in ISO/IEC 9075-3.

### Function

Specify the cursor of a <preparable dynamic delete statement: positioned> or a <preparable dynamic update statement: positioned>.

### Format

```
<preparable dynamic cursor name> ::=
 [<scope option>] <cursor name>
```

### Syntax Rules

- 1) Let *PDCN* be the <preparable dynamic cursor name>, let *CN* be the <cursor name> contained in *PDCN*, and let *P* be the <prepare statement>, <preparable dynamic delete statement: positioned>, or <preparable dynamic update statement: positioned> that contains *PDCN*.

Case:

- a) If *PDCN* contains a <scope option> that specifies GLOBAL, then

Case:

- i) If there exists in the current SQL-session context an extended dynamic cursor *EDC* with an <extended cursor name> having a global scope and a <cursor name> that is equivalent to *CN*, then *EDC* is the cursor referenced by *PDCN*.

- ii) Otherwise, *PDCN* is said to be *invalid*.

- b) If *PDCN* contains a <scope option> that specifies LOCAL, or if no <scope option> is specified, then:

- i) 03 The *potentially referenced cursors* of *PDCN* include every declared dynamic cursor whose <cursor name> is equivalent to *CN* and whose scope is the containing SQL-client module (minus any <SQL schema statement>s contained in the SQL-client module) and every extended dynamic cursor having an <extended cursor name> that has a scope of the containing SLI-client module (minus any <SQL schema statement>s contained in the SQL-client module) and whose <cursor name> is equivalent to *CN*.

- ii) Case:

- 1) If the number of potentially referenced cursors is greater than 1 (one), then *PDCN* is said to be *ambiguous*.
- 2) If the number of potentially referenced cursors is less than 1 (one), then *PDCN* is said to be *invalid*.
- 3) Otherwise, *PDCN* refers to the single potentially referenced cursor of *P*.

**Access Rules**

*None.*

**General Rules**

*None.*

**Conformance Rules**

*None.*

## 20.24 <preparable dynamic update statement: positioned>

### Function

Update a row of a table through a dynamic cursor.

### Format

```
<preparable dynamic update statement: positioned> ::=
 UPDATE [<target table>] SET <set clause list>
 WHERE CURRENT OF <preparable dynamic cursor name>
```

### Syntax Rules

- 1) If the <preparable dynamic cursor name> *PDCN* is not ambiguous or invalid, then:
  - a) Let *CR* be the cursor identified by *PDCN*. *CR* shall be an updatable cursor.
  - b) If <target table> is not specified, then let *QE* be the <query expression> simply contained in the <cursor specification> of the result set descriptor of the cursor identified by <preparable dynamic cursor name>. Let *LUT* be the leaf underlying table of *QE* such that *QE* is one-to-one with respect to *LUT*. Let *TN* be the name of *LUT*.
 

Case:

    - i) If the <table reference> that references *LUT* specifies ONLY, then the <target table>
 

ONLY ( *TN* )

is implicit.
    - ii) Otherwise, the <target table>
 

*TN*

is implicit.
  - c) Let *TT* be the explicit or implicit <target table>.
  - d) Let *TU* be the simply underlying table of *CR*. *TU* is the *subject table* of the <preparable dynamic update statement: positioned> *PDUSP*. Let *LUT* be the leaf underlying table of *TU* such that *TU* is one-to-one with respect to *LUT*.
  - e) Let *TN* be the <table name> contained in *TT*. *TN* shall identify *LUT*.
  - f) *TT* shall specify ONLY if and only if the <table reference> contained in *TU* that references *LUT* specifies ONLY.
  - g) *TN* shall not identify an old transition table or a new transition table.
  - h) Let *T* be the table identified by *TN*.
  - i) Let *CN* be the <table name> contained in <target table>. *CN* is an exposed <table or query name>.

**20.24 <preparable dynamic update statement: positioned>**

- j) The scope of *CN* is <set clause list>.
- k) If *CR* is an ordered cursor, then for each <object column> *OC* contained in <set clause list>, the <order by clause> that is simply contained in the <cursor specification> for *CR* shall not generally contain a <column reference> that references *OC* or an underlying column of the column identified by *OC*.
- l) Each <column name> specified as an <object column> shall identify a column in the explicit or implicit <column name list> contained in the explicit or implicit <updatability clause> of the <cursor specification> included in the result set descriptor of *CR*.

**Access Rules**

- 1) If *PDCN* is not ambiguous or invalid, then

Case:

- a) If <preparable dynamic update statement: positioned> is contained, without an intervening <SQL routine spec> that specifies SQL SECURITY INVOKER, in an <SQL schema statement>, then let *A* be the <authorization identifier> that owns that schema. The applicable privileges for *A* shall include UPDATE for each <object column>.
- b) Otherwise, the current privileges shall include UPDATE for each <object column>.

**General Rules**

- 1) Let *SCL* be the <set clause list>.
- 2) The General Rules of Subclause 15.6, “Effect of a positioned update”, are applied with *CR* as *CURSOR*, *SCL* as *SET CLAUSE LIST*, *PDUSP* as *STATEMENT*, and *TT* as *TARGET*

NOTE 462 — If the General Rules are reached, then *PDCN* cannot be ambiguous or invalid.

**Conformance Rules**

- 1) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <preparable dynamic update statement: positioned>.

*(Blank page)*



## 21 Embedded SQL

### 21.1 <embedded SQL host program>

*This Subclause is modified by Subclause 16.1, “<embedded SQL host program>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 8.1, “<embedded SQL host program>”, in ISO/IEC 9075-10.  
This Subclause is modified by Subclause 18.1, “<embedded SQL host program>”, in ISO/IEC 9075-14.*

#### Function

Specify an <embedded SQL host program>.

#### Format

```
[10] <embedded SQL host program> ::=
 <embedded SQL Ada program>
 | <embedded SQL C program>
 | <embedded SQL COBOL program>
 | <embedded SQL Fortran program>
 | <embedded SQL MUMPS program>
 | <embedded SQL Pascal program>
 | <embedded SQL PL/I program>

<embedded SQL statement> ::=
 <SQL prefix> <statement or declaration> [<SQL terminator>]

[10] <statement or declaration> ::=
 <declare cursor>
 | <dynamic declare cursor>
 | <temporary table declaration>
 | <embedded authorization declaration>
 | <embedded path specification>
 | <embedded transform group specification>
 | <embedded collation specification>
 | <embedded exception declaration>
 | <SQL procedure statement>

[10] <SQL prefix> ::=
 EXEC SQL
 | <ampersand>SQL<left paren>

<SQL terminator> ::=
 END-EXEC
 | <semicolon>
 | <right paren>

<embedded authorization declaration> ::=
 DECLARE <embedded authorization clause>
```

## 21.1 &lt;embedded SQL host program&gt;

```

<embedded authorization clause> ::=
 SCHEMA <schema name>
 | AUTHORIZATION <embedded authorization identifier>
 [FOR STATIC { ONLY | AND DYNAMIC }]
 | SCHEMA <schema name> AUTHORIZATION <embedded authorization identifier>
 [FOR STATIC { ONLY | AND DYNAMIC }]

<embedded authorization identifier> ::=
 <module authorization identifier>

<embedded path specification> ::=
 <path specification>

<embedded transform group specification> ::=
 <transform group specification>

<embedded collation specification> ::=
 <module collations>

<embedded SQL declare section> ::=
 <embedded SQL begin declare>
 [<embedded character set declaration>]
 [<host variable definition>...]
 <embedded SQL end declare>
 | <embedded SQL MUMPS declare>

<embedded character set declaration> ::=
 SQL NAMES ARE <character set specification>

<embedded SQL begin declare> ::=
 <SQL prefix> BEGIN DECLARE SECTION [<SQL terminator>]

<embedded SQL end declare> ::=
 <SQL prefix> END DECLARE SECTION [<SQL terminator>]

<embedded SQL MUMPS declare> ::=
 <SQL prefix>
 BEGIN DECLARE SECTION
 [<embedded character set declaration>]
 [<host variable definition>...]
 END DECLARE SECTION
 <SQL terminator>

<host variable definition> ::=
 <Ada variable definition>
 | <C variable definition>
 | <COBOL variable definition>
 | <Fortran variable definition>
 | <MUMPS variable definition>
 | <Pascal variable definition>
 | <PL/I variable definition>

10 <embedded variable name> ::=
 <colon><host identifier>

<host identifier> ::=
 <Ada host identifier>
 | <C host identifier>
 | <COBOL host identifier>

```

| <Fortran host identifier>  
| <MUMPS host identifier>  
| <Pascal host identifier>  
| <PL/I host identifier>

## Syntax Rules

- 1) An <embedded SQL host program> is a compilation unit that consists of programming language text and SQL text. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s, as defined in this International Standard. The programming language text shall conform to the requirements of a specific programming language, the *host language*. When <embedded SQL Ada program>, <embedded SQL C program>, <embedded SQL COBOL program>, <embedded SQL Fortran program>, <embedded SQL MUMPS program>, <embedded SQL Pascal program>, or <embedded SQL PL/I program> is specified, the host language is Ada, C, COBOL, Fortran, M, Pascal, or PL/I, respectively.

NOTE 463 — “Compilation unit” is defined in Subclause 4.22, “SQL-client modules”.

- 2) 10 Case:
- a) An <embedded SQL statement> or <embedded SQL MUMPS declare> that is contained in an <embedded SQL MUMPS program> shall contain an <SQL prefix> that is “<ampersand>SQL<left paren>”. There shall be no <separator> between the <ampersand> and “SQL” nor between “SQL” and the <left paren>.
  - b) An <embedded SQL statement>, <embedded SQL begin declare>, or <embedded SQL end declare> that is not contained in an <embedded SQL MUMPS program> shall contain an <SQL prefix> that is “EXEC SQL”.
- 3) Case:
- a) An <embedded SQL statement>, <embedded SQL begin declare>, or <embedded SQL end declare> contained in an <embedded SQL COBOL program> shall contain an <SQL terminator> that is END-EXEC.
  - b) An <embedded SQL statement>, <embedded SQL begin declare>, or <embedded SQL end declare> contained in an <embedded SQL Fortran program> shall not contain an <SQL terminator>.
  - c) 10 An <embedded SQL statement>, <embedded SQL begin declare>, or <embedded SQL end declare> contained in an <embedded SQL Ada program>, <embedded SQL C program>, <embedded SQL Pascal program>, or <embedded SQL PL/I program> shall contain an <SQL terminator> that is a <semicolon>.
  - d) An <embedded SQL statement> or <embedded SQL MUMPS declare> that is contained in an <embedded SQL MUMPS program> shall contain an <SQL terminator> that is a <right paren>.
- 4) Case:
- a) An <embedded SQL declare section> that is contained in an <embedded SQL MUMPS program> shall be an <embedded SQL MUMPS declare>.
  - b) An <embedded SQL declare section> that is not contained in an <embedded SQL MUMPS program> shall not be an <embedded SQL MUMPS declare>.

10 NOTE 464 — There is no restriction on the number of <embedded SQL declare section>s that may be contained in an <embedded SQL host program>.

## 21.1 &lt;embedded SQL host program&gt;

- 5) The <token>s comprising an <SQL prefix>, <embedded SQL begin declare>, or <embedded SQL end declare> shall be separated by <space> characters and shall be specified on one line. Otherwise, the rules for the continuation of lines and tokens from one line to the next and for the placement of host language comments are those of the programming language of the containing <embedded SQL host program>.
- 6) If an <embedded authorization declaration> appears in an <embedded SQL host program>, then it shall be contained in the first <embedded SQL statement> of that <embedded SQL host program>.
- 7) An <embedded SQL host program> shall not contain more than one <embedded path specification>.
- 8) An <embedded SQL host program> shall not contain more than one <embedded transform group specification>.
- 9) An <embedded SQL host program> shall not contain more than one <embedded collation specification>.
- 10) Case:
  - a) If <embedded transform group specification> is not specified, then an <embedded transform group specification> containing a <multiple group specification> with a <group specification> *GS* for each <host variable definition> that has an associated user-defined type *UDT*, but is not a user-defined locator variable is implicit. The <group name> of *GS* is implementation-defined and its <path-resolved user-defined type name> is the <user-defined type name> of *UDT*.
  - b) If <embedded transform group specification> contains a <single group specification> with a <group name> *GN*, then an <embedded transform group specification> containing a <multiple group specification> with a <group specification> *GS* for each <host variable definition> that has an associated user-defined type *UDT*, but is not a user-defined type locator variable is implicit. The <group name> of *GS* is *GN* and its <path-resolved user-defined type name> is the <user-defined type name> of *UDT*.
  - c) If <embedded transform group specification> contains a <multiple group specification> *MGS*, then an <embedded transform group specification> containing a <multiple group specification> that contains *MGS* extended with a <group specification> *GS* for each <host variable definition> that has an associated user-defined type *UDT*, but is not a user-defined locator variable and no equivalent of *UDT* is contained in any <group specification> contained in *MGS* is implicit. The <group name> of *GS* is implementation-defined and its <path-resolved user-defined type name> is the <user-defined type name> of *UDT*.
- 11) In the text of the <embedded SQL host program>, the implicit or explicit <embedded transform group specification> shall precede every <host variable definition>.
- 12) An <embedded SQL host program> shall contain no more than one <embedded character set declaration>. If an <embedded character set declaration> is not specified, then an <embedded character set declaration> that specifies an implementation-defined character set that contains at least every character that is in <SQL language character> is implicit.
- 13) A <temporary table declaration> that is contained in an <embedded SQL host program> shall precede in the text of that <embedded SQL host program> any SQL-statement or <declare cursor> that references the <table name> of the <temporary table declaration>.
- 14) A <declare cursor> that is contained in an <embedded SQL host program> shall precede in the text of that <embedded SQL host program> any SQL-statement that references the <cursor name> of the <declare cursor>.
- 15) A <dynamic declare cursor> that is contained in an <embedded SQL host program> shall precede in the text of that <embedded SQL host program> any SQL-statement that references the <cursor name> of the <dynamic declare cursor>.

- 16) 04 10 Any <host identifier> that is contained in an <embedded SQL statement> in an <embedded SQL host program> shall be defined in exactly one <host variable definition> contained in that <embedded SQL host program>. In programming languages that support <host variable definition>s in subprograms, two <host variable definition>s with different, non-overlapping scope in the host language are to be regarded as defining different host variables, even if they specify the same variable name. That <host variable definition> shall appear in the text of the <embedded SQL host program> prior to any <embedded SQL statement> that references the <host identifier>. The <host variable definition> shall be such that a host language reference to the <host identifier> is valid at every <embedded SQL statement> that contains the <host identifier>.
- 17) 10 A <host variable definition> defines the host language data type of the <host identifier>. For every such host language data type, when the host language is Ada, C, COBOL, Fortran, M, Pascal, or PL/I, an equivalent SQL <data type> is specified in Subclause 21.3, “<embedded SQL Ada program>”, Subclause 21.4, “<embedded SQL C program>”, Subclause 21.5, “<embedded SQL COBOL program>”, Subclause 21.6, “<embedded SQL Fortran program>”, Subclause 21.7, “<embedded SQL MUMPS program>”, Subclause 21.8, “<embedded SQL Pascal program>”, and Subclause 21.9, “<embedded SQL PL/I program>”, respectively.
- 18) 10 An <embedded SQL host program> shall contain a <host variable definition> that specifies SQLSTATE.
- 19) If one or more <host variable definition>s that specify SQLSTATE appear in an <embedded SQL host program>, then the <host variable definition>s shall be such that a host language reference to SQLSTATE is valid at every <embedded SQL statement>, including <embedded SQL statement>s that appear in any subprograms contained in that <embedded SQL host program>. The first such <host variable definition> of SQLSTATE shall appear in the text of the <embedded SQL host program> prior to any <embedded SQL statement>.
- 20) 10 Given an <embedded SQL host program> *H*, there is an implied standard-conforming <SQL-client module definition> *M* and an implied host program *P* derived from *H*. The derivation of the implied program *P* and the implied <SQL-client module definition> *M* of an <embedded SQL host program> *H* effectively precedes the processing of any host language program text manipulation commands such as inclusion or copying of text.

NOTE 465 — Before *H* can be executed, *M* is processed by an implementation-defined mechanism to produce an SQL-client module. An SQL-implementation may combine this mechanism with the processing of the <embedded SQL host program>, in which the existence of *M* is pure hypothetical.

Given an <embedded SQL host program> *H* with an implied <SQL-client module definition> *M* and an implied program *P* defined as above:

- a) The implied <SQL-client module definition> *M* of *H* shall be a standard-conforming <SQL-client module definition>.
  - b) The implied program *P* shall conform to the specification of the host language.
- 21) 10 *M* is derived from *H* as follows:
- a) *M* contains a <module name clause> whose <SQL-client module name> is either implementation-dependent or is omitted.
  - b) *M* contains a <module character set specification> that is identical to the explicit or implicit <embedded character set declaration> with the keyword “SQL” removed.
  - c) *M* contains a <language clause> that specifies either ADA, C, COBOL, FORTRAN, M, PASCAL, or PLI, where *H* is respectively an <embedded SQL Ada program>, an <embedded SQL C program>, an

## 21.1 &lt;embedded SQL host program&gt;

<embedded SQL COBOL program>, an <embedded SQL Fortran program>, an <embedded SQL MUMPS program>, an <embedded SQL Pascal program>, or an <embedded SQL PL/I program>.

d) Case:

- i) If *H* contains an <embedded authorization declaration> *EAD*, then let *EAC* be the <embedded authorization clause> contained in *EAD*; *M* contains a <module authorization clause> that specifies *EAC*.
- ii) Otherwise, let *SN* be an implementation-defined <schema name>; *M* contains a <module authorization clause> that specifies “SCHEMA *SN*”.

e) Case:

- i) If *H* contains an <embedded path specification> *EPS*, then *M* contains the <module path specification> *EPS*.
- ii) Otherwise, *M* contains an implementation-defined <module path specification>.
- f) *M* contains a <module transform group specification> that is identical to the explicit or implicit <embedded transform group specification>.
- g) If an <embedded collation specification> *ECS* is specified, then *M* contains a <module collations> that is identical to the <module collations> contained in *ECS*.
- h) For every <declare cursor> *EC* contained in *H*, *M* contains one <declare cursor> *PC* and one <externally-invoked procedure> *PS* that contains an <open statement> that references *PC*.
  - i) The <procedure name> of *PS* is implementation-dependent. *PS* contains a <host parameter declaration> *PD* for each distinct <embedded variable name> *EVN* contained in *PC* with an implementation-dependent <host parameter name> *PN* and the <host parameter data type> *PT*, determined as follows.

Case:

- 1) If *EVN* identifies a binary large object locator variable, then *PT* is BLOB AS LOCATOR.
- 2) If *EVN* identifies a character large object locator variable, then *PT* is CLOB AS LOCATOR.
- 3) If *EVN* identifies an array locator variable, then *PT* is AAT AS LOCATOR, where *AAT* is the associated array type of *V*.
- 4) If *EVN* identifies a multiset locator variable, then *PT* is AMT AS LOCATOR, where *AMT* is the associated multiset type of *V*.
- 5) 14 If *EVN* identifies a user-defined type locator variable, then *PT* is UDT AS LOCATOR, where *UDT* is the associated user-defined type of *V*.
- 6) Otherwise, *PT* is the SQL data type that corresponds to the host language data type of *EVN* as specified in Subclause 13.6, “Data type correspondences”.
- ii) *PS* contains a <host parameter declaration> that specifies SQLSTATE. The order of <host parameter declaration>s in *PS* is implementation-dependent. *PC* is a copy of *EC* in which each *EVN* has been replaced as follows.

Case:

- 1) If *EVN* does not identify user-defined type locator variable, but *EVN* identifies a host variable that has an associated user-defined type *UT*, then:
  - A) Let *GN* be the <group name> corresponding to the <user-defined type name> of *UT* contained in <group specification> contained in <embedded transform group specification>.
  - B) Apply the Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, with *DT* and *GN* as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function *TSF*.
  - C) Let the declared type of the single SQL parameter of *TSF* be *TPT*. *PT* shall be assignable to *TPT*.
  - D) *EVN* is replaced by:

*TSFN*(*CAST*(*PN* AS *TPT*))

- 2) 14 Otherwise, *EVN* is replaced by:

*PN*

- i) For every <dynamic declare cursor> *EC* in *H*, *M* contains one <dynamic declare cursor> *PC* that is a copy of *EC*.
- j) *M* contains one <temporary table declaration> for each <temporary table declaration> contained in *H*. Each <temporary table declaration> of *M* is a copy of the corresponding <temporary table declaration> of *H*.
- k) 04 *M* contains one <embedded exception declaration> for each <embedded exception declaration> contained in *H*. Each <embedded exception declaration> of *M* is a copy of the corresponding <embedded exception declaration> of *H*.
- l) 04 *M* contains an <externally-invoked procedure> for each <SQL procedure statement> contained in *H*. The <externally-invoked procedure> *PS* of *M* corresponding with an <SQL procedure statement> *ES* of *H* is defined as follows:

Case:

- i) If *ES* is not an <open statement>, then:
  - 1) The <procedure name> of *PS* is implementation-dependent.
  - 2) Let *n* be the number of distinct <embedded variable name>s contained in *ES*. Let *HVN<sub>i</sub>*, 1 (one) ≤ *i* ≤ *n*, be the *i*-th such <embedded variable name> and let *HV<sub>i</sub>* be the host variable identified by *HVN<sub>i</sub>*.
  - 3) For each *HVN<sub>i</sub>*, 1 (one) ≤ *i* ≤ *n*, *PS* contains a <host parameter declaration> *PD<sub>i</sub>* defining a host parameter *P<sub>i</sub>* such that:
    - A) The <host parameter name> *PN<sub>i</sub>* of *PD<sub>i</sub>* is implementation-dependent.
    - B) The <host parameter data type> *PT<sub>i</sub>* of *PD<sub>i</sub>* is determined as follows.

Case:

- I) If  $HV_i$  is a binary large object locator variable, then  $PT_i$  is BLOB AS LOCATOR.
  - II) If  $HV_i$  is a character large object locator variable, then  $PT_i$  is CLOB AS LOCATOR.
  - III) If  $HV_i$  is an array locator variable, then  $PT_i$  is AAT AS LOCATOR, where AAT is the associated array type of  $HV_i$ .
  - IV) If  $HV_i$  is a multiset locator variable, then  $PT_i$  is AMT AS LOCATOR, where AMT is the associated multiset type of  $HV_i$ .
  - V) 14 If  $HV_i$  is user-defined type locator variable, then  $PT_i$  is UDT AS LOCATOR, where UDT is the associated user-defined type of  $HV_i$ .
  - VI) Otherwise,  $PT_i$  is the SQL data type that corresponds to the host language data type of  $HV_i$  as specified in Subclause 13.6, “Data type correspondences”.
- 4)  $PS$  contains a <host parameter declaration> that specifies SQLSTATE.
- 5) The order of the <host parameter declaration>s  $PD_i$ ,  $1 \text{ (one)} \leq i \leq n$ , is implementation-dependent.
- 6) For each  $HVN_i$ ,  $1 \text{ (one)} \leq i \leq n$ , that identifies some  $HV_i$  that has an associated user-defined type, but is not a user-defined type locator variable, the Syntax Rules of Subclause 9.6, “Host parameter mode determination”, are applied with  $PD_i$  corresponding to  $HVN_i$  as *HOST PARAM DECL* and  $ES$  as *SQL PROC STMT* to determine whether the corresponding  $P_i$  is an input host parameter, an output host parameter, or both an input host parameter and an output host parameter.
- A) Among  $P_i$ ,  $1 \text{ (one)} \leq i \leq n$ , let  $a$  be the number of input host parameters,  $b$  be the number of output host parameters, and let  $c$  be the number of host parameters that are both input host parameters and output host parameters.
  - B) Among  $P_i$ ,  $1 \text{ (one)} \leq i \leq n$ , let  $PI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the input host parameters, let  $PO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the output host parameters, and let  $PIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the host parameters that are both input host parameters and output host parameters.
  - C) Let  $PNI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the <host parameter name> of  $PI_j$ . Let  $PNO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the <host parameter name> of  $PO_k$ . Let  $PNIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the <host parameter name> of  $PIO_l$ .
  - D) Let  $HVI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the host variable corresponding to  $PI_j$ . Let  $HVO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the host variable corresponding to  $PO_k$ . Let  $HVIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the host variable corresponding to  $PIO_l$ .
  - E) Let  $TSI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the associated SQL data type of  $HVI_j$ . Let  $TSO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the associated SQL data type of  $HVO_k$ . Let  $TSIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the associated SQL data type of  $HVIO_l$ .



- F) Let  $TUI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the associated user-defined type of  $HVI_j$ . Let  $TUO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the associated user-defined type of  $HVO_k$ . Let  $TUIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the associated user-defined type of  $HVIO_l$ .
- G) Let  $GNI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the <group name> corresponding to the <user-defined type name> of  $TUI_j$  contained in the <group specification> contained in <embedded transform group specification>. Let  $GNO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the <group name> corresponding to the <user-defined type name> of  $TUO_k$  contained in the <group specification> contained in <embedded transform group specification>. Let  $GNIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the <group name> corresponding to the <user-defined type name> of  $TUIO_l$  contained in the <group specification> contained in <embedded transform group specification>.
- H) For every  $j$ ,  $1 \text{ (one)} \leq j \leq a$ , apply the Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, with  $TUI_j$  and  $GNI_j$  as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function  $TSFI_j$  identified by <routine name>  $TSIN_j$ . Let  $TTI_j$  be the data type of the single SQL parameter of  $TSFI_j$ .  $TSI_j$  shall be assignable to  $TTI_j$ .
- I) For every  $l$ ,  $1 \text{ (one)} \leq l \leq c$ , apply the Syntax Rules of Subclause 9.21, “Determination of a to-sql function”, with  $TUIO_l$  and  $GNIO_l$  as *TYPE* and *GROUP*, respectively. There shall be an applicable to-sql function  $TSFIO_l$  identified by <routine name>  $TSION_l$ . Let  $TTIO_l$  be the data type of the single SQL parameter of  $TSFIO_l$ .  $TSIO_l$  shall be assignable to  $TTIO_l$ .
- J) For every  $k$ ,  $1 \text{ (one)} \leq k \leq b$ , apply the Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, with  $TUO_k$  and  $GNO_k$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSFO_k$  identified by <routine name>  $FSON_k$ . Let  $TRO_k$  be the result data type of  $FSFO_k$ .  $TSO_k$  shall be assignable to  $TRO_k$ .
- K) For every  $l$ ,  $1 \text{ (one)} \leq l \leq c$ , apply the Syntax Rules of Subclause 9.19, “Determination of a from-sql function”, with  $TUIO_l$  and  $GNIO_l$  as *TYPE* and *GROUP*, respectively. There shall be an applicable from-sql function  $FSFIO_l$  identified by <routine name>  $FSION_l$ . Let  $TRIO_l$  be the result data type of  $FSFIO_l$ .  $TSIO_l$  shall be assignable to  $TRIO_l$ .
- L) Let  $SVI_j$ ,  $1 \text{ (one)} \leq j \leq a$ ,  $SVO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , and  $SVIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be implementation-dependent <SQL variable name>s, each of which is not equivalent to any other <SQL variable name> contained in *ES*, to any <SQL parameter name> contained in *ES*, or to any <column name> contained in *ES*.
- 7) 14 Let *NES* be an <SQL procedure statement> that is a copy of *ES* in which every  $HVN_i$ ,  $1 \text{ (one)} \leq i \leq n$ , is replaced as follows:
- Case:
- A) If  $HV_i$  has an associated user-defined type but is not a user-defined type locator variable, then
- Case:

## 21.1 &lt;embedded SQL host program&gt;

- I) If  $P_i$  is an input host parameter, then let  $PI_j$ ,  $1 \text{ (one)} \leq j \leq a$ , be the input host parameter that corresponds to  $P_i$ ;  $HVN_i$  is replaced by  $SVI_j$ .
  - II) If  $P_i$  is an output host parameter, then let  $PO_k$ ,  $1 \text{ (one)} \leq k \leq b$ , be the output host parameter that corresponds to  $P_i$ ;  $HVN_i$  is replaced by  $SVO_k$ .
  - III) Otherwise, let  $PIO_l$ ,  $1 \text{ (one)} \leq l \leq c$ , be the input host parameter and the output host parameter that corresponds to  $P_i$ ;  $HVN_i$  is replaced by  $SVIO_l$ .
- B) 14 Otherwise,  $HVN_i$  is replaced by  $PN_i$ .
- 8) 14 The <SQL procedure statement> of  $PS$  is:

```

BEGIN ATOMIC
 DECLARE SVI_1 TUI_1 ;
 ...
 DECLARE SVI_a TUI_a ;
 DECLARE SVO_1 TUO_1 ;
 ...
 DECLARE SVO_b TUO_b ;
 DECLARE $SVIO_1$ $TUIO_1$;
 ...
 DECLARE $SVIO_c$ $TUIO_c$;
 SET $SVI_1 = TSIN_1$ (CAST (PNI_1 AS TTI_1));
 ...
 SET $SVI_a = TSIN_a$ (CAST (PNI_a AS TTI_a));
 SET $SVIO_1 = TSION_1$ (CAST ($PNIO_1$ AS $TTIO_1$));
 ...
 SET $SVIO_c = TSION_c$ (CAST ($PNIO_c$ AS $TTIO_c$));
 NES;
 SET $PNO_1 = CAST$ ($FSON_1$ (SVO_1) AS TSO_1);
 ...
 SET $PNO_b = CAST$ ($FSON_b$ (SVO_b) AS TSO_b);
 SET $PNIO_1 = CAST$ ($FSION_1$ ($SVIO_1$) AS $TSIO_1$);
 ...
 SET $PNIO_c = CAST$ ($FSION_c$ ($SVIO_c$) AS $TSIO_c$);
END;
```

- 9) Whether one <externally-invoked procedure> of  $M$  can correspond to more than one <SQL procedure statement> of  $H$  is implementation-dependent.
- ii) If  $ES$  is an <open statement>, then:
- 1) Let  $EC$  be the <declare cursor> in  $H$  referenced by  $ES$ .
  - 2)  $PS$  is the <externally-invoked procedure> in  $M$  that contains an <open statement> that references the <declare cursor> in  $M$  corresponding to  $EC$ .
- 22) 10  $P$  is derived from  $H$  as follows:
- a) Each <embedded SQL begin declare>, <embedded SQL end declare>, and <embedded character set declaration> has been deleted. If the embedded host language is M, then each <embedded SQL MUMPS declare> has been deleted.

- b) Each <host variable definition> in an <embedded SQL declare section> has been replaced by a valid data definition in the target host language according to the Syntax Rules specified in an <embedded SQL Ada program>, <embedded SQL C program>, <embedded SQL COBOL program>, <embedded SQL Fortran program>, <embedded SQL Pascal program>, or an <embedded SQL PL/I program> clause.
- c) 04 Each <embedded SQL statement> that contains a <declare cursor>, a <dynamic declare cursor>, an <SQL-invoked routine>, or a <temporary table declaration> has been deleted, and every <embedded SQL statement> that contains an <embedded exception declaration> has been replaced with statements of the host language that will have the effect specified by the General Rules of Subclause 21.2, “<embedded exception declaration>”.
- d) Each <embedded SQL statement> that contains an <SQL procedure statement> has been replaced by host language statements that perform the following actions:
  - i) A host language procedure or subroutine call of the <externally-invoked procedure> of the implied <SQL-client module definition> *M* of *H* that corresponds with the <SQL procedure statement>.

If the <SQL procedure statement> is not an <open statement>, then the arguments of the call include each distinct <host identifier> contained in the <SQL procedure statement> together with the SQLSTATE <host identifier>. If the <SQL procedure statement> is an <open statement>, then the arguments of the call include each distinct <host identifier> contained in the corresponding <declare cursor> of *H* together with the SQLSTATE <host identifier>.

The order of the arguments in the call corresponds with the order of the corresponding <host parameter declaration>s in the corresponding <externally-invoked procedure>.

NOTE 466 — In an <embedded SQL Fortran program>, the “SQLSTATE” variable may be abbreviated to “SQLSTA”. See the Syntax Rules of Subclause 21.6, “<embedded SQL Fortran program>”.
  - ii) Exception actions, as specified in Subclause 21.2, “<embedded exception declaration>”.
- e) Each <statement or declaration> that contains an <embedded authorization declaration> is deleted.

## Access Rules

- 1) For every host variable whose <embedded variable name> is contained in <statement or declaration> and has an associated user-defined type, the current privileges shall include EXECUTE privilege on all from-sql functions (if any) and all to-sql functions (if any) referenced in the corresponding SQL-client module.

## General Rules

- 1) The interpretation of an <embedded SQL host program> *H* is defined to be equivalent to the interpretation of the implied program *P* of *H* and the implied <SQL-client module definition> *M* of *H*.

## Conformance Rules

- 1) Without Feature B051, “Enhanced execution rights”, conforming SQL language shall not contain an <embedded authorization declaration>.
- 2) Without Feature F461, “Named character sets”, conforming SQL language shall not contain an <embedded character set declaration>.

**21.1 <embedded SQL host program>**

- 3) Without Feature F361, “Subprogram support”, conforming SQL language shall not contain two <host variable definition>s that specify the same variable name.
- 4) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain an <embedded path specification>.
- 5) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <embedded transform group specification>.

## 21.2 <embedded exception declaration>

### Function

Specify the action to be taken when an SQL-statement causes a specific class of condition to be raised.

### Format

```
<embedded exception declaration> ::=
 WHENEVER <condition> <condition action>

<condition> ::=
 <SQL condition>

<SQL condition> ::=
 <major category>
 | SQLSTATE (<SQLSTATE class value> [, <SQLSTATE subclass value>])
 | CONSTRAINT <constraint name>

<major category> ::=
 SQLEXCEPTION
 | SQLWARNING
 | NOT FOUND

<SQLSTATE class value> ::=
 <SQLSTATE char><SQLSTATE char> !! See the Syntax Rules.

<SQLSTATE subclass value> ::=
 <SQLSTATE char><SQLSTATE char><SQLSTATE char> !! See the Syntax Rules.

<SQLSTATE char> ::=
 <simple Latin upper case letter>
 | <digit>

<condition action> ::=
 CONTINUE
 | <go to>

<go to> ::=
 { GOTO | GO TO } <goto target>

<goto target> ::=
 <host label identifier>
 | <unsigned integer>
 | <host PL/I label variable>

<host label identifier> ::=
 !! See the Syntax Rules.

<host PL/I label variable> ::=
 !! See the Syntax Rules.
```

## Syntax Rules

- 1) SQLWARNING, NOT FOUND, and SQLEXCEPTION correspond to SQLSTATE class values corresponding to categories W, N, and X in Table 33, “SQLSTATE class and subclass values”, respectively.
- 2) An <embedded exception declaration> contained in an <embedded SQL host program> applies to an <SQL procedure statement> contained in that <embedded SQL host program> if and only if the <SQL procedure statement> appears after the <embedded exception declaration> that has condition *C* in the text sequence of the <embedded SQL host program> and no other <embedded exception declaration> *E* that satisfies one of the following conditions appears between the <embedded exception declaration> and the <SQL procedure statement> in the text sequence of the <embedded SQL host program>.

Let *D* be the <condition> contained in *E*.

- a) *D* is the same as *C*.
  - b) *D* is a <major category> and belongs to the same class to which *C* belongs.
  - c) *D* contains an <SQLSTATE class value>, but does not contain an <SQLSTATE subclass value>, and *E* contains the same <SQLSTATE class value> that *C* contains.
  - d) *D* contains the <SQLSTATE class value> that corresponds to *integrity constraint violation* and *C* contains CONSTRAINT.
- 3) In the values of <SQLSTATE class value> and <SQLSTATE subclass value>, there shall be no <separator> between the <SQLSTATE char>s.
  - 4) The values of <SQLSTATE class value> and <SQLSTATE subclass value> shall correspond to class values and subclass values, respectively, specified in Table 33, “SQLSTATE class and subclass values”.
  - 5) If an <embedded exception declaration> specifies a <go to>, then the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> shall be such that a host language GO TO statement specifying that <host label identifier>, <host PL/I label variable>, or <unsigned integer> is valid at every <SQL procedure statement> to which the <embedded exception declaration> applies.

NOTE 467 —

If an <embedded exception declaration> is contained in an <embedded SQL Ada program>, then the <goto target> of a <go to> should specify a <host label identifier> that is a label\_name in the containing <embedded SQL Ada program>.

If an <embedded exception declaration> is contained in an <embedded SQL C program>, then the <goto target> of a <go to> should specify a <host label identifier> that is a label in the containing <embedded SQL C program>.

If an <embedded exception declaration> is contained in an <embedded SQL COBOL program>, then the <goto target> of a <go to> should specify a <host label identifier> that is a section-name or an unqualified paragraph-name in the containing <embedded SQL COBOL program>.

If an <embedded exception declaration> is contained in an <embedded SQL Fortran program>, then the <goto target> of a <go to> should be an <unsigned integer> that is the statement label of an executable statement that appears in the same program unit as the <go to>.

If an <embedded exception declaration> is contained in an <embedded SQL MUMPS program>, then the <goto target> of a <go to> should be a gotoargument that is the statement label of an executable statement that appears in the same <embedded SQL MUMPS program>.

If an <embedded exception declaration> is contained in an <embedded SQL Pascal program>, then the <goto target> of a <go to> should be an <unsigned integer> that is a label.

If an <embedded exception declaration> is contained in an <embedded SQL PL/I program>, then the <goto target> of a <go to> should specify either a <host label identifier> or a <host PL/I label variable>.

Case:

- If <host label identifier> is specified, then the <host label identifier> should be a label constant in the containing <embedded SQL PL/I program>.
- If <host PL/I label variable> is specified, then the <host PL/I label variable> should be a PL/I label variable declared in the containing <embedded SQL PL/I program>.

## Access Rules

*None.*

## General Rules

- 1) Immediately after the execution of an <SQL procedure statement> *STMT* in an <embedded SQL host program> that returns an SQLSTATE value other than *successful completion*:
  - a) Let *E* be the set of <embedded exception declaration>s that are contained in the <embedded SQL host program> containing *STMT*, that applies to *STMT*, and that specifies a <condition action> that is <go to>.
  - b) Let *CV* and *SCV* be respectively the values of the class and subclass of the SQLSTATE value that indicates the result of the <SQL procedure statement>.
  - c) If the execution of the <SQL procedure statement> caused the violation of one or more constraints or assertions, then:
    - i) Let *ECN* be the set of <embedded exception declaration>s in *E* that specify CONSTRAINT and the <constraint name> of a constraint that was violated by execution of *STMT*.
    - ii) If *ECN* contains more than one <embedded exception declaration>, then an implementation-dependent <embedded exception declaration> is chosen from *ECN*; otherwise, the single <embedded exception declaration> in *ECN* is chosen.
    - iii) A GO TO statement of the host language is performed, specifying the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> specified in the <embedded exception declaration> chosen from *ECN*.
  - d) Otherwise:
    - i) Let *ECS* be the set of <embedded exception declaration>s in *E* that specify SQLSTATE, an <SQLSTATE class value>, and an <SQLSTATE subclass value>.
    - ii) If *ECS* contains an <embedded exception declaration> *EY* that specifies an <SQLSTATE class value> identical to *CV* and an <SQLSTATE subclass value> identical to *SCV*, then a GO TO statement of the host language is performed, specifying the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> specified in the <embedded exception declaration> *EY*.
    - iii) Otherwise:
      - 1) Let *EC* be the set of <embedded exception declaration>s in *E* that specify SQLSTATE and an <SQLSTATE class value> without an <SQLSTATE subclass value>.

## 21.2 &lt;embedded exception declaration&gt;

- 2) If *EC* contains an <embedded exception declaration> *EY* that specifies an <SQLSTATE class value> identical to *CV*, then a GO TO statement of the host language is performed, specifying the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> specified in the <embedded exception declaration> *EY*.
- 3) Otherwise:
  - A) Let *EX* be the set of <embedded exception declaration>s in *E* that specify SQLEXCEPTION.
  - B) If *EX* contains an <embedded exception declaration> *EY* and *CV* belongs to Category X in Table 33, “SQLSTATE class and subclass values”, then a GO TO statement of the host language is performed, specifying the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> specified in the <embedded exception declaration> *EY*.
  - C) Otherwise:
    - I) Let *EW* be the set of <embedded exception declaration>s in *E* that specify SQLWARNING.
    - II) If *EW* contains an <embedded exception declaration> *EY* and *CV* belongs to Category W in Table 33, “SQLSTATE class and subclass values”, then a GO TO statement of the host language is performed, specifying the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> specified in the <embedded exception declaration> *EY*.
    - III) Otherwise, let *ENF* be the set of <embedded exception declaration>s in *E* that specify NOT FOUND. If *ENF* contains an <embedded exception declaration> *EY* and *CV* belongs to Category N in Table 33, “SQLSTATE class and subclass values”, then a GO TO statement of the host language is performed, specifying the <host label identifier>, <host PL/I label variable>, or <unsigned integer> of the <go to> specified in the <embedded exception declaration> *EY*.

## Conformance Rules

- 1) Without Feature B041, “Extensions to embedded SQL exception declarations”, conforming SQL language shall not contain an <SQL condition> that contains either SQLSTATE or CONSTRAINT.
- 2) Without Feature F491, “Constraint management”, conforming SQL language shall not contain an <SQL condition> that contains a <constraint name>.



## 21.3 <embedded SQL Ada program>

*This Subclause is modified by Subclause 18.1, “<embedded SQL Ada program>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 18.2, “<embedded SQL Ada program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL Ada program>.

### Format

```
<embedded SQL Ada program> ::=
 !! See the Syntax Rules.

<Ada variable definition> ::=
 <Ada host identifier> [{ <comma> <Ada host identifier> }...] <colon>
 <Ada type specification> [<Ada initial value>]

<Ada initial value> ::=
 <Ada assignment operator> <character representation>...

<Ada assignment operator> ::=
 <colon><equals operator>

<Ada host identifier> ::=
 !! See the Syntax Rules.

<Ada type specification> ::=
 <Ada qualified type specification>
 | <Ada unqualified type specification>
 | <Ada derived type specification>

<Ada qualified type specification> ::=
 Interfaces.SQL.CHAR
 [CHARACTER SET [IS] <character set specification>]
 <left paren> 1 <double period> <character length> <right paren>
 | Interfaces.SQL.SMALLINT
 | Interfaces.SQL.INT
 | Interfaces.SQL.BIGINT
 | Interfaces.SQL.REAL
 | Interfaces.SQL.DOUBLE_PRECISION
 | Interfaces.SQL.BOOLEAN
 | Interfaces.SQL.SQLSTATE_TYPE
 | Interfaces.SQL.INDICATOR_TYPE

<Ada unqualified type specification> ::=
 CHAR <left paren> 1 <double period> <character length> <right paren>
 | SMALLINT
 | INT
 | BIGINT
 | REAL
 | DOUBLE_PRECISION
 | BOOLEAN
 | SQLSTATE_TYPE
 | INDICATOR_TYPE
```

```

09 14 <Ada derived type specification> ::=
 <Ada CLOB variable>
 | <Ada CLOB locator variable>
 | <Ada BINARY variable>
 | <Ada VARBINARY variable>
 | <Ada BLOB variable>
 | <Ada BLOB locator variable>
 | <Ada user-defined type variable>
 | <Ada user-defined type locator variable>
 | <Ada REF variable>
 | <Ada array locator variable>
 | <Ada multiset locator variable>

<Ada CLOB variable> ::=
 SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<Ada CLOB locator variable> ::=
 SQL TYPE IS CLOB AS LOCATOR

<Ada BINARY variable> ::=
 SQL TYPE IS BINARY <left paren> <length> <right paren>

<Ada VARBINARY variable> ::=
 SQL TYPE IS VARBINARY <left paren> <length> <right paren>

<Ada BLOB variable> ::=
 SQL TYPE IS BLOB <left paren> <large object length> <right paren>

<Ada BLOB locator variable> ::=
 SQL TYPE IS BLOB AS LOCATOR

<Ada user-defined type variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS <predefined type>

<Ada user-defined type locator variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

<Ada REF variable> ::=
 SQL TYPE IS <reference type>

<Ada array locator variable> ::=
 SQL TYPE IS <array type> AS LOCATOR

<Ada multiset locator variable> ::=
 SQL TYPE IS <multiset type> AS LOCATOR

```

## Syntax Rules

- 1) An <embedded SQL Ada program> is a compilation unit that consists of Ada text and SQL text. The Ada text shall conform to [ISO8652]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.
- 2) An <embedded SQL statement> may be specified wherever an Ada statement may be specified. An <embedded SQL statement> may be prefixed by an Ada label.
- 3) An <Ada host identifier> is any valid Ada identifier. An <Ada host identifier> shall be contained in an <embedded SQL Ada program>.

- 4) An <Ada variable definition> defines one or more host variables.
- 5) An <Ada variable definition> shall be modified as follows before it is placed into the program derived from the <embedded SQL Ada program> (see the Syntax Rules of [Subclause 21.1](#), “<embedded SQL host program>”):

- a) Any optional CHARACTER SET specification shall be removed from an <Ada qualified type specification> and <Ada derived type specification>.

Let  $k$  be

Case:

- i) If OCTETS is specified, then 1 (one).
- ii) If a <character set specification> is specified, then the maximum number of octets per character of the character set identified by the <character set specification>.
- iii) Otherwise, the maximum number of octets per character of the implementation-defined default character set.
- b) Any optional <char length units> shall be removed from an <Ada qualified type specification> and <Ada derived type specification>.
- c) The <character length>  $CL$  specified in a CHAR declaration of any <Ada qualified type specification> shall be replaced by a length whose value is equal to  $L*k$ , where  $L$  is the value of the <length> contained in  $CL$ .
- d) The syntax

```
SQL TYPE IS BINARY (L)
```

for a given <Ada host identifier>  $HVN$  shall be replaced by

```
HVN : Interfaces.SQL.CHAR(1..L)
```

in any <Ada BINARY variable>.

- e) The syntax

```
SQL TYPE IS VARBINARY (L)
```

for a given <Ada host identifier>  $HVN$  shall be replaced by

```
TYPE HVN IS RECORD
 HVN_RESERVED : Interfaces.SQL.INT;
 HVN_LENGTH : Interfaces.SQL.INT;
 HVN_DATA : Interfaces.SQL.CHAR(1..L);
END RECORD;
```

in any <Ada VARBINARY variable>.

- f) 14 The syntax

```
SQL TYPE IS CLOB (L)
```

for a given <Ada host identifier>  $HVN$  shall be replaced by

```

TYPE HVN IS RECORD
 HVN_RESERVED : Interfaces.SQL.INT;
 HVN_LENGTH : Interfaces.SQL.INT;
 HVN_DATA : Interfaces.SQL.CHAR(1..LL);
END RECORD;

```

in any <Ada CLOB variable>, where the value of *LL* is *NV*\**k*, and *NV* is the value of *L* as specified in Subclause 6.1, “<data type>”.

g) The syntax

```
SQL TYPE IS BLOB (L)
```

for a given <Ada host identifier> *HVN* shall be replaced by

```

TYPE HVN IS RECORD
 HVN_RESERVED : Interfaces.SQL.INT;
 HVN_LENGTH : Interfaces.SQL.INT;
 HVN_DATA : Interfaces.SQL.CHAR(1..L);
END RECORD;

```

in any <Ada BLOB variable>.

h) The syntax

```
SQL TYPE IS UDTN AS PDT
```

shall be replaced by

```
ADT
```

in any <Ada user-defined type variable>, where *ADT* is the data type listed in the “Ada data type” column corresponding to the row for SQL data type *PDT* in Table 16, “Data type correspondences for Ada”. *ADT* shall not be “none”. The data type identified by *UDTN* is called the *associated user-defined type* of the host variable and the data type identified by *PDT* is called the *associated SQL data type* of the host variable.

i) The syntax

```
SQL TYPE IS BLOB AS LOCATOR
```

shall be replaced by

```
Interfaces.SQL.INT
```

in any <Ada BLOB locator variable>. The host variable defined by <Ada BLOB locator variable> is called a *binary large object locator variable*.

j) The syntax

```
SQL TYPE IS CLOB AS LOCATOR
```

shall be replaced by

Interfaces.SQL.INT

in any <Ada CLOB locator variable>. The host variable defined by <Ada CLOB locator variable> is called a *character large object locator variable*.

k) The syntax

SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

shall be replaced by

Interfaces.SQL.INT

in any <Ada user-defined type locator variable>. The host variable defined by <Ada user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

l) The syntax

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

Interfaces.SQL.INT

in any <Ada array locator variable>. The host variable defined by <Ada array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

m) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

Interfaces.SQL.INT

in any <Ada multiset locator variable>. The host variable defined by <Ada multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

n) 09 The syntax

SQL TYPE IS <reference type>

for a given <Ada host identifier> *RTV* shall be replaced by

*RTV* : Interfaces.SQL.CHAR(1..*RL*)

in any <Ada REF variable>, where the value of *RL* is the length in octets of the reference type.

The modified <Ada variable definition> shall be a valid Ada object-declaration in the program derived from the <embedded SQL Ada program>.

- 6) The reference type identified by <reference type> contained in an <Ada REF variable> is called the *referenced type* of the reference.

**21.3 <embedded SQL Ada program>**

- 7) An <Ada variable definition> shall be specified within the scope of Ada **with** and **use** clauses that specify the following:

```
with Interfaces.SQL;
use Interfaces.SQL;
use Interfaces.SQL.CHARACTER_SET;
```

- 8) The <character representation> sequence in an <Ada initial value> specifies an initial value to be assigned to the Ada variable. It shall be a valid Ada specification of an initial value.
- 9) CHAR describes a character string variable whose equivalent SQL data type is CHARACTER with the same length and character set specified by <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit.
- 10) SMALLINT, INT, and BIGINT describe exact numeric variables. The equivalent SQL data types are SMALLINT, INTEGER, and BIGINT, respectively.
- 11) REAL and DOUBLE\_PRECISION describe approximate numeric variables. The equivalent SQL data types are REAL and DOUBLE PRECISION, respectively.
- 12) BOOLEAN describes a boolean variable. The equivalent SQL data type is BOOLEAN.
- 13) SQLSTATE\_TYPE describes a character string variable whose length is the length of the SQLSTATE parameter, five characters.
- 14) INDICATOR\_TYPE describes an exact numeric variable whose specific data type is any <exact numeric type> with a scale of 0 (zero).

**Access Rules**

*None.*

**General Rules**

- 1) See Subclause 21.1, “<embedded SQL host program>”.

**Conformance Rules**

- 1) Without Feature B011, “Embedded Ada”, conforming SQL language shall not contain an <embedded SQL Ada program>.
- 2) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada BLOB variable>.
- 3) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada CLOB variable>.
- 4) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada BLOB locator variable>.
- 5) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada CLOB locator variable>.

- 6) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain an <Ada qualified type specification> that contains Interfaces.SQL.BIGINT.
- 7) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain an <Ada unqualified type specification> that contains BIGINT.
- 8) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain an <Ada BINARY variable>.
- 9) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain an <Ada VARBINARY variable>.
- 10) Without Feature S241, “Transform functions”, conforming SQL language shall not contain an <Ada user-defined type variable>.
- 11) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain an <Ada REF variable>.
- 12) Without Feature S232, “Array locators”, conforming SQL language shall not contain an <Ada array locator variable>.
- 13) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain an <Ada multiset locator variable>.
- 14) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in an <Ada user-defined type locator variable> that identifies a structured type.

## 21.4 <embedded SQL C program>

*This Subclause is modified by Subclause 18.2, “<embedded SQL C program>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 18.3, “<embedded SQL C program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL C program>.

### Format

```
<embedded SQL C program> ::=
 !! See the Syntax Rules.

<C variable definition> ::=
 [<C storage class>] [<C class modifier>]
 <C variable specification> <semicolon>

<C variable specification> ::=
 <C numeric variable>
 | <C character variable>
 | <C derived variable>

<C storage class> ::=
 auto
 | extern
 | static

<C class modifier> ::=
 const
 | volatile

<C numeric variable> ::=
 { long long | long | short | float | double }
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [<C initial value>] }...]

<C character variable> ::=
 <C character type> [CHARACTER SET [IS] <character set specification>]
 <C host identifier> <C array specification> [<C initial value>]
 [{ <comma> <C host identifier> <C array specification>
 [<C initial value>] }...]

<C character type> ::=
 char
 | unsigned char
 | unsigned short

<C array specification> ::=
 <left bracket> <character length> <right bracket>

<C host identifier> ::=
 !! See the Syntax Rules.

09 14 <C derived variable> ::=
```



```

 <C VARCHAR variable>
| <C NCHAR variable>
| <C NCHAR VARYING variable>
| <C CLOB variable>
| <C NCLOB variable>
| <C BINARY variable>
| <C VARBINARY variable>
| <C BLOB variable>
| <C user-defined type variable>
| <C CLOB locator variable>
| <C BLOB locator variable>
| <C array locator variable>
| <C multiset locator variable>
| <C user-defined type locator variable>
| <C REF variable>

<C VARCHAR variable> ::=
 VARCHAR [CHARACTER SET [IS] <character set specification>]
 <C host identifier> <C array specification> [<C initial value>]
 [{ <comma> <C host identifier> <C array specification> [
 <C initial value>] }...]

<C NCHAR variable> ::=
 NCHAR [CHARACTER SET [IS] <character set specification>]
 <C host identifier> <C array specification> [<C initial value>]
 [{ <comma> <C host identifier> <C array specification>
 [<C initial value>] } ...]

<C NCHAR VARYING variable> ::=
 NCHAR VARYING [CHARACTER SET [IS] <character set specification>]
 <C host identifier> <C array specification> [<C initial value>]
 [{ <comma> <C host identifier> <C array specification> [
 <C initial value>] } ...]

<C CLOB variable> ::=
 SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]
 <C host identifier> [<C initial value>] [{ <comma> <C host identifier> [
 <C initial value>] }...]

<C NCLOB variable> ::=
 SQL TYPE IS NCLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]
 <C host identifier> [<C initial value>] [{ <comma> <C host identifier>
 [<C initial value>] }...]

<C user-defined type variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS <predefined type>
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] } ...]

<C BINARY variable> ::=
 SQL TYPE IS BINARY <left paren> <length> <right paren>
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [<C initial value>] }...]

<C VARBINARY variable> ::=
 SQL TYPE IS VARBINARY <left paren> <length> <right paren>

```

**21.4 <embedded SQL C program>**

```

 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [<C initial value>] }...]

<C BLOB variable> ::=
 SQL TYPE IS BLOB <left paren> <large object length> <right paren>
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] } ...]

<C CLOB locator variable> ::=
 SQL TYPE IS CLOB AS LOCATOR
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] } ...]

<C BLOB locator variable> ::=
 SQL TYPE IS BLOB AS LOCATOR
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] } ...]

<C array locator variable> ::=
 SQL TYPE IS <array type> AS LOCATOR
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] } ...]

<C multiset locator variable> ::=
 SQL TYPE IS <multiset type> AS LOCATOR
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] } ...]

<C user-defined type locator variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR
 <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [
 <C initial value>] }...]

<C REF variable> ::=
 SQL TYPE IS <reference type> <C host identifier> [<C initial value>]
 [{ <comma> <C host identifier> [<C initial value>] }...]

<C initial value> ::=
 <equals operator> <character representation>...

```

**Syntax Rules**

- 1) An <embedded SQL C program> is a compilation unit that consists of C text and SQL text. The C text shall conform to [ISO9899]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.
- 2) An <embedded SQL statement> may be specified wherever a C statement may be specified within a function block. If the C statement could include a label prefix, then the <embedded SQL statement> may be immediately preceded by a label prefix.

- 3) A <C host identifier> is any valid C variable identifier. A <C host identifier> shall be contained in an <embedded SQL C program>.
- 4) A <C variable definition> defines one or more host variables.
- 5) A <C variable definition> shall be modified as follows before it is placed into the program derived from the <embedded SQL C program> (see the Syntax Rules of [Subclause 21.1](#), “<embedded SQL host program>”):

- a) 14 Any optional CHARACTER SET specification shall be removed from a <C VARCHAR variable>, a <C character variable>, a <C CLOB variable>, a <C NCHAR variable>, <C NCHAR VARYING variable>, or a <C NCLOB variable>.

Let  $k$  be

Case:

- i) If OCTETS is specified, then 1 (one).
- ii) If a <character set specification> is specified, then the maximum number of octets per character of the character set identified by the <character set specification>.
- iii) Otherwise, the maximum number of octets per character of the implementation-defined default character set.
- b) Any optional <char length units> shall be removed from a <character length> or a <character large object length>.
- c) The syntax “VARCHAR” shall be replaced by “char” in any <C VARCHAR variable>.
- d) The <character length>  $CL$  specified in a <C array specification> in any <C character variable> whose <C character type> specifies “char” or “unsigned char”, in any <C VARCHAR variable>, in any <C NCHAR variable>, or in any <C NCHAR VARYING variable> shall be replaced by a length whose value is equal to  $NV * k$ , where  $NV$  is the numeric value of the <length> contained in  $CL$ .

NOTE 468 — The <character length> is not adjusted for <C character type>s that specify “unsigned short” because the units of <character length> are presumed to already be the same units as used by the underlying character set.

- e) The syntax “NCHAR” in any <C NCHAR variable> and the syntax “NCHAR VARYING” in any <C NCHAR VARYING variable> shall be replaced by “char”.
- f) 14 The syntax

SQL TYPE IS CLOB (  $L$  )

and the syntax

SQL TYPE IS NCLOB (  $L$  )

for a given <C host identifier>  $hvn$  shall be replaced by

```
struct {
 long hvn_reserved;
 unsigned long hvn_length;
```

## 21.4 &lt;embedded SQL C program&gt;

```

char hvn_data[LL];
} hvn

```

in any <C CLOB variable> or <C NCLOB variable>, where the value of *LL* is  $NV \cdot k$  and *NV* is the numeric value of *L* as defined in Subclause 6.1, “<data type>”.

## g) The syntax

```
SQL TYPE IS BLOB (L)
```

for a given <C host identifier> *hvn* shall be replaced by:

```

struct {
 long hvn_reserved;
 unsigned long hvn_length;
 char hvn_data[L];
} hvn

```

in any <C BLOB variable>.

## h) The syntax

```
SQL TYPE IS BINARY (L)
```

for a given <C host identifier> *hvn* shall be replaced by

```
unsigned char HVN [L]
```

in any <C BINARY variable>.

## i) The syntax

```
SQL TYPE IS VARBINARY (L)
```

for a given <C host identifier> *hvn* shall be replaced by

```

struct {
 long hvn_reserved;
 unsigned long hvn_length;
 char hvn_data [L];
} hvn

```

in any <C VARBINARY variable>.

## j) The syntax

```
SQL TYPE IS UDTN AS PDT
```

shall be replaced by

```
ADT
```

in any <C user-defined type variable>, where *ADT* is the data type listed in the “C data type” column corresponding to the row for SQL data type *PDT* in Table 17, “Data type correspondences for C”. *ADT* shall not be “none”. The data type identified by *UDTN* is called the *associated user-defined type* of the host variable and the data type identified by *PDT* is called the *associated SQL data type* of the host variable.

k) The syntax

SQL TYPE IS BLOB AS LOCATOR

shall be replaced by

unsigned long

in any <C BLOB locator variable>. The host variable defined by <C BLOB locator variable> is called a *binary large object locator variable*.

l) The syntax

SQL TYPE IS CLOB AS LOCATOR

shall be replaced by

unsigned long

in any <C CLOB locator variable>. The host variable defined by <C CLOB locator variable> is called a *character large object locator variable*.

m) The syntax

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

unsigned long

in any <C array locator variable>. The host variable defined by <C array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

n) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

unsigned long

in any <C multiset locator variable>. The host variable defined by <C multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

o) The syntax

SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

shall be replaced by

unsigned long

in any <C user-defined type locator variable>. The host variable defined by <C user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

p) 09 The syntax

```
SQL TYPE IS <reference type>
```

for a given <C host identifier> *hvn* shall be replaced by

```
unsigned char hvn[L]
```

in any <C REF variable>, where the value of *L* is the length in octets of the reference type.

The modified <C variable definition> shall be a valid C data declaration in the program derived from the <embedded SQL C program>.

- 6) The reference type identified by <reference type> contained in a <C REF variable> is called the *referenced type* of the reference.
- 7) The <character representation> sequence contained in a <C initial value> specifies an initial value to be assigned to the C variable. It shall be a valid C specification of an initial value.
- 8) Except for array specifications for character strings, a <C variable definition> shall specify a scalar type.
- 9) 14 In a <C variable definition>, the words “VARCHAR”, “CHARACTER”, “SET”, “IS”, “VARYING”, “BINARY”, “VARBINARY”, “BLOB”, “CLOB”, “NCHAR”, “NCLOB”, “AS”, “LOCATOR”, and “REF” may be specified in any combination of upper-case and lower-case letters (see the Syntax Rules of Subclause 5.2, “<token> and <separator>”).
- 10) 14 In a <C character variable>, a <C VARCHAR variable>, or a <C CLOB variable>, if a <character set specification> is specified, then the equivalent SQL data type is CHARACTER, CHARACTER VARYING, or CHARACTER LARGE OBJECT whose character set is the same as the character set specified by the <character set specification>. In a <C NCHAR variable>, a <C NCHAR VARYING variable>, or a <C NCLOB variable>, if a <character set specification> is specified, then the equivalent SQL data type is NATIONAL CHARACTER, NATIONAL CHARACTER VARYING, or NATIONAL CHARACTER LARGE OBJECT whose character set is the same as the character set specified by the <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit.
- 11) Each <C host identifier> specified in a <C character variable> or a <C NCHAR variable> describes a fixed-length character string. The length is specified by the <character length> of the <C array specification>, measured in the units specified by the implicit or explicit <char length units>. The value in the host variable is terminated by a null character and the position occupied by this null character is included in the length of the host variable. The equivalent SQL data type is CHARACTER or NATIONAL CHARACTER, respectively, whose length is one less than the <character length> of the <C array specification> and whose value does not include the terminating null character. The <character length> shall be greater than 1 (one).
- 12) Each <C host identifier> specified in a <C VARCHAR variable> or a <C NCHAR VARYING variable> describes a variable-length character string. The maximum length is specified by the <character length> of the <C array specification>, measured in the units specified by the implicit or explicit <char length units>. The value in the host variable is terminated by a null character and the position occupied by this null character is included in the maximum length of the host variable. The equivalent SQL data type is CHARACTER VARYING or NATIONAL CHARACTER VARYING, respectively, whose maximum length is 1 (one) less than the <character length> of the <C array specification> and whose value does not include the terminating null character. The <character length> shall be greater than 1 (one).
- 13) “short” describes an exact numeric variable. The equivalent SQL data type is SMALLINT.

- 14) “long” describes an exact numeric variable. The equivalent SQL data type is INTEGER or BOOLEAN.
- 15) “long long” describes an exact numeric variable. The equivalent SQL data type is BIGINT.
- 16) “float” describes an approximate numeric variable. The equivalent SQL data type is REAL.
- 17) “double” describes an approximate numeric variable. The equivalent SQL data type is DOUBLE PRECISION.

## Access Rules

*None.*

## General Rules

- 1) See Subclause 21.1, “<embedded SQL host program>”.

## Conformance Rules

- 1) Without Feature B012, “Embedded C”, conforming SQL language shall not contain an <embedded SQL C program>.
- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <C REF variable>.
- 3) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <C user-defined type variable>.
- 4) Without Feature S232, “Array locators”, conforming SQL language shall not contain an <C array locator variable>.
- 5) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <C multiset locator variable>.
- 6) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <C user-defined type locator variable> that identifies a structured type.
- 7) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C BLOB variable>.
- 8) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C CLOB variable>.
- 9) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C BLOB locator variable>.
- 10) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C CLOB locator variable>.
- 11) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain a <C numeric variable> that contains long long.

- 12) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <C BINARY variable>.
- 13) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <C VARBINARY variable>.



## 21.5 <embedded SQL COBOL program>

*This Subclause is modified by Subclause 18.3, “<embedded SQL COBOL program>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 18.4, “<embedded SQL COBOL program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL COBOL program>.

### Format

```
<embedded SQL COBOL program> ::=
 !! See the Syntax Rules.

<COBOL variable definition> ::=
 { 01 | 77 } <COBOL host identifier>
 <COBOL type specification> [<character representation>...] <period>

<COBOL host identifier> ::=
 !! See the Syntax Rules.

<COBOL type specification> ::=
 <COBOL character type>
 | <COBOL national character type>
 | <COBOL numeric type>
 | <COBOL integer type>
 | <COBOL derived type specification>

09 14 <COBOL derived type specification> ::=
 <COBOL CLOB variable>
 | <COBOL NCLOB variable>
 | <COBOL BINARY variable>
 | <COBOL BLOB variable>
 | <COBOL user-defined type variable>
 | <COBOL CLOB locator variable>
 | <COBOL BLOB locator variable>
 | <COBOL array locator variable>
 | <COBOL multiset locator variable>
 | <COBOL user-defined type locator variable>
 | <COBOL REF variable>

<COBOL character type> ::=
 [CHARACTER SET [IS] <character set specification>]
 { PIC | PICTURE } [IS] { X [<left paren> <character length> <right paren>] }...

<COBOL national character type> ::=
 [CHARACTER SET [IS] <character set specification>]
 { PIC | PICTURE } [IS] { N [<left paren> <character length> <right paren>] }...

<COBOL CLOB variable> ::=
 [USAGE [IS]] SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<COBOL NCLOB variable> ::=
```

**21.5 <embedded SQL COBOL program>**

```

 [USAGE [IS]] SQL TYPE IS NCLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<COBOL BINARY variable> ::=
 [USAGE [IS]] SQL TYPE IS BINARY <left paren> <length> <right paren>

<COBOL BLOB variable> ::=
 [USAGE [IS]] SQL TYPE IS BLOB <left paren> <large object length> <right paren>

<COBOL user-defined type variable> ::=
 [USAGE [IS]] SQL TYPE IS <path-resolved user-defined type name>
 AS <predefined type>

<COBOL CLOB locator variable> ::=
 [USAGE [IS]] SQL TYPE IS CLOB AS LOCATOR

<COBOL BLOB locator variable> ::=
 [USAGE [IS]] SQL TYPE IS BLOB AS LOCATOR

<COBOL array locator variable> ::=
 [USAGE [IS]] SQL TYPE IS <array type> AS LOCATOR

<COBOL multiset locator variable> ::=
 [USAGE [IS]] SQL TYPE IS <multiset type> AS LOCATOR

<COBOL user-defined type locator variable> ::=
 [USAGE [IS]] SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

<COBOL REF variable> ::=
 [USAGE [IS]] SQL TYPE IS <reference type>

<COBOL numeric type> ::=
 { PIC | PICTURE } [IS] S <COBOL nines specification>
 [USAGE [IS]] DISPLAY SIGN LEADING SEPARATE

<COBOL nines specification> ::=
 <COBOL nines> [V [<COBOL nines>]]
 | V <COBOL nines>

<COBOL integer type> ::=
 <COBOL binary integer>

<COBOL binary integer> ::=
 { PIC | PICTURE } [IS] S<COBOL nines>
 [USAGE [IS]] BINARY

<COBOL nines> ::=
 { 9 [<left paren> <length> <right paren>] }...

```

NOTE 469 — The syntax “N(L)” is not part of the current COBOL standard, so its use is merely a recommendation; therefore, the production <COBOL national character type> is not normative in this edition of ISO/IEC 9075.

**Syntax Rules**

- 1) An <embedded SQL COBOL program> is a compilation unit that consists of COBOL text and SQL text. The COBOL text shall conform to [ISO1989]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.

- 2) An <embedded SQL statement> in an <embedded SQL COBOL program> may be specified wherever a COBOL statement may be specified in the Procedure Division of the <embedded SQL COBOL program>. If the COBOL statement could be immediately preceded by a paragraph-name, then the <embedded SQL statement> may be immediately preceded by a paragraph-name.
- 3) A <COBOL host identifier> is any valid COBOL data-name. A <COBOL host identifier> shall be contained in an <embedded SQL COBOL program>.
- 4) A <COBOL variable definition> is a restricted form of COBOL data description entry that defines a host variable.
- 5) A <COBOL variable definition> shall be modified as follows before it is placed into the program derived from the <embedded SQL COBOL program> (see the Syntax Rules of [Subclause 21.1](#), “<embedded SQL host program>”).
  - a) Any optional CHARACTER SET specification shall be removed from a <COBOL character type>, a <COBOL national character type>, a <COBOL CLOB variable>, and a <COBOL NCLOB variable>.

Let  $k$  be

Case:

- i) If OCTETS is specified, then 1 (one).
  - ii) If a <character set specification> is specified, then the maximum number of octets per character of the character set identified by the <character set specification>.
  - iii) Otherwise, the maximum number of octets per character of the implementation-defined default character set.
- b) Any optional <char length units> shall be removed from a <character length> or a <character large object length>.
  - c) The <character length>  $CL$  specified in any <COBOL character type> shall be replaced by a length equal to  $NV \cdot k$ , where  $NV$  is the numeric value of the <length> contained in  $CL$ .

NOTE 470 — The <character length> specified in a <COBOL national character type> is not adjusted, because the units of <character length> are presumed to already be the same units as used by the underlying character set.
  - d) 14 The syntax

SQL TYPE IS CLOB (  $L$  )

or the syntax

SQL TYPE IS NCLOB (  $L$  )

for a given <COBOL host identifier>  $HVN$  shall be replaced by:

```
49 HVN-RESERVED PIC S9(9) USAGE IS BINARY.
49 HVN-LENGTH PIC S9(9) USAGE IS BINARY.
49 HVN-DATA PIC X(LL).
```

in any <COBOL CLOB variable> or <COBOL NCLOB variable>, where the value of  $LL$  is  $NV \cdot k$  and  $NV$  is the numeric value of  $L$  as defined in [Subclause 6.1](#), “<data type>”.

- e) The syntax

SQL TYPE IS BLOB ( L )

for a given <COBOL host identifier> *HVN* shall be replaced by:

```
49 HVN-RESERVED PIC S9(9) USAGE IS BINARY.
49 HVN-LENGTH PIC S9(9) USAGE IS BINARY.
49 HVN-DATA PIC X(LL).
```

in any <COBOL BLOB variable>.

f) The syntax

SQL TYPE IS BINARY ( L )

shall be replaced by

PICTURE X (L)

in any <COBOL BINARY variable>.

g) The syntax

SQL TYPE IS *UDTN* AS *PDT*

shall be replaced by

*ADT*

in any <COBOL user-defined type variable>, where *ADT* is the data type listed in the “COBOL data type” column corresponding to the row for SQL data type *PDT* in [Table 18, “Data type correspondences for COBOL”](#). *ADT* shall not be “none”. The data type identified by *UDTN* is called the *associated user-defined type* of the host variable and the data type identified by *PDT* is called the *associated SQL data type* of the host variable.

h) The syntax

SQL TYPE IS BLOB AS LOCATOR

shall be replaced by

PIC S9(9) USAGE IS BINARY

in any <COBOL BLOB locator variable>. The host variable defined by <COBOL BLOB locator variable> is called a *binary large object locator variable*.

i) The syntax

SQL TYPE IS CLOB AS LOCATOR

shall be replaced by

PIC S9(9) USAGE IS BINARY

in any <COBOL CLOB locator variable>. The host variable defined by <COBOL CLOB locator variable> is called a *character large object locator variable*.

j) The syntax

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

PIC S9(9) USAGE IS BINARY

in any <COBOL array locator variable>. The host variable defined by <COBOL array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

k) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

PIC S9(9) USAGE IS BINARY

in any <COBOL multiset locator variable>. The host variable defined by <COBOL multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

l) The syntax

SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

shall be replaced by

PIC S9(9) USAGE IS BINARY

in any <COBOL user-defined type locator variable>. The host variable defined by <COBOL user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

m) 09 The syntax

SQL TYPE IS <reference type>

for a given <COBOL host identifier> *HVN* shall be replaced by

01 *HVN* PICTURE X(*RL*)

in any <COBOL REF variable>, where the value of *RL* is the length in octets of the reference type.

The modified <COBOL variable definition> shall be a valid data description entry in the Data Division of the program derived from the <embedded SQL COBOL program>.

- 6) The reference type identified by <reference type> contained in a <COBOL REF variable> is called the *referenced type* of the reference.
- 7) The optional <character representation> sequence in a <COBOL variable definition> may specify a VALUE clause. Whether other clauses may be specified is implementation-defined. The <character representation> sequence shall be such that the <COBOL variable definition> is a valid COBOL data description entry.

- 8) 14 A <COBOL character type> describes a character string variable whose equivalent SQL data type is CHARACTER with the same length and character set specified by <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit.
- 9) A <COBOL numeric type> describes an exact numeric variable. The equivalent SQL data type is NUMERIC of the same precision and scale.
- 10) A <COBOL binary integer> describes an exact numeric variable. The equivalent SQL data type is SMALLINT, INTEGER, or BIGINT.

## Access Rules

*None.*

## General Rules

- 1) See Subclause 21.1, “<embedded SQL host program>”.

## Conformance Rules

- 1) Without Feature B013, “Embedded COBOL”, conforming SQL language shall not contain an <embedded SQL COBOL program>.
- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <COBOL REF variable>.
- 3) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <COBOL user-defined type variable>.
- 4) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <COBOL array locator variable>.
- 5) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <COBOL multiset locator variable>.
- 6) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <COBOL user-defined type locator variable> that identifies a structured type.
- 7) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL BLOB variable>.
- 8) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL CLOB variable>.
- 9) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL BLOB locator variable>.
- 10) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL CLOB locator variable>.

- 11) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <COBOL BINARY variable>.

## 21.6 <embedded SQL Fortran program>

*This Subclause is modified by Subclause 18.4, “<embedded SQL Fortran program>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 18.5, “<embedded SQL Fortran program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL Fortran program>.

### Format

```

<embedded SQL Fortran program> ::=
 !! See the Syntax Rules.

<Fortran variable definition> ::=
 <Fortran type specification> <Fortran host identifier>
 [{ <comma> <Fortran host identifier> }...]

<Fortran host identifier> ::=
 !! See the Syntax Rules.

<Fortran type specification> ::=
 CHARACTER [<asterisk> <character length>] [CHARACTER SET
 [IS] <character set specification>]
 | CHARACTER KIND = n [<asterisk> <character length>]
 [CHARACTER SET [IS] <character set specification>]
 | INTEGER
 | REAL
 | DOUBLE PRECISION
 | LOGICAL
 | <Fortran derived type specification>

09 14 <Fortran derived type specification> ::=
 <Fortran CLOB variable>
 | <Fortran BINARY variable>
 | <Fortran VARBINARY variable>
 | <Fortran BLOB variable>
 | <Fortran user-defined type variable>
 | <Fortran CLOB locator variable>
 | <Fortran BLOB locator variable>
 | <Fortran user-defined type locator variable>
 | <Fortran array locator variable>
 | <Fortran multiset locator variable>
 | <Fortran REF variable>

<Fortran CLOB variable> ::=
 SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<Fortran BINARY variable> ::=
 SQL TYPE IS BINARY <left paren> <length> <right paren>

<Fortran VARBINARY variable> ::=
 SQL TYPE IS VARBINARY <left paren> <length> <right paren>

```



```
<Fortran BLOB variable> ::=
 SQL TYPE IS BLOB <left paren> <large object length> <right paren>

<Fortran user-defined type variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS <predefined type>

<Fortran CLOB locator variable> ::=
 SQL TYPE IS CLOB AS LOCATOR

<Fortran BLOB locator variable> ::=
 SQL TYPE IS BLOB AS LOCATOR

<Fortran user-defined type locator variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

<Fortran array locator variable> ::=
 SQL TYPE IS <array type> AS LOCATOR

<Fortran multiset locator variable> ::=
 SQL TYPE IS <multiset type> AS LOCATOR

<Fortran REF variable> ::=
 SQL TYPE IS <reference type>
```

## Syntax Rules

- 1) An <embedded SQL Fortran program> is a compilation unit that consists of Fortran text and SQL text. The Fortran text shall conform to [ISO1539]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.
- 2) An <embedded SQL statement> may be specified wherever an executable Fortran statement may be specified. An <embedded SQL statement> that precedes any executable Fortran statement in the containing <embedded SQL Fortran program> shall not have a Fortran statement number. Otherwise, if the Fortran statement could have a statement number then the <embedded SQL statement> can have a statement number.
- 3) Blanks are significant in <embedded SQL statement>s. The rules for <separator>s in an <embedded SQL statement> are as specified in [Subclause 5.2, “<token> and <separator>”](#).
- 4) A <Fortran host identifier> is any valid Fortran variable name with all <space> characters removed. A <Fortran host identifier> shall be contained in an <embedded SQL Fortran program>.
- 5) A <Fortran variable definition> is a restricted form of Fortran type-statement that defines one or more host variables.
- 6) A <Fortran variable definition> shall be modified as follows before it is placed into the program derived from the <embedded SQL Fortran program> (see the Syntax Rules [Subclause 21.1, “<embedded SQL host program>”](#)).
  - a) Any optional CHARACTER SET specification shall be removed from the CHARACTER and the CHARACTER KIND=*n* alternatives in a <Fortran type specification>.  
Let *k* be  
Case:
    - i) If OCTETS is specified, then 1 (one).

## 21.6 &lt;embedded SQL Fortran program&gt;

- ii) If a <character set specification> is specified, then the maximum number of octets per character of the character set identified by the <character set specification>.
  - iii) Otherwise, the maximum number of octets per character of the implementation-defined default character set.
- b) Any optional <char length units> shall be removed from a <character length> or a <character large object length>.
- c) The <character length> *CL* specified in the CHARACTER alternative of any <Fortran type specification> shall be replaced by a length equal to  $NV \cdot k$ , where *NV* is the numeric value of the <length> contained in *CL*.

NOTE 471 — The <character length> is not adjusted for CHARACTER KIND=*n* alternatives of any <Fortran type specification>, because the units of <character length> are presumed to already be the same units used by the underlying character set.

- d) 14 The syntax

```
SQL TYPE IS CLOB (L)
```

for a given <Fortran host identifier> *HVN* shall be replaced by

```
INTEGER HVN_RESERVED
INTEGER HVN_LENGTH
CHARACTER HVN_DATA [* LL]
```

in any <Fortran CLOB variable>, where the value of *LL* is  $NV \cdot k$  and *NV* is the numeric value of *L* as specified in [Subclause 6.1](#), “<data type>”.

- e) The syntax

```
SQL TYPE IS BLOB (L)
```

for a given <Fortran host identifier> *HVN* shall be replaced by

```
INTEGER HVN_RESERVED
INTEGER HVN_LENGTH
CHARACTER HVN_DATA [* L]
```

in any <Fortran BLOB variable>.

- f) The syntax

```
SQL TYPE IS BINARY (L)
```

for a given <Fortran host identifier> *HVN* shall be replaced by

```
CHARACTER HVN * L
```

in any <Fortran BINARY variable>.

- g) The syntax

```
SQL TYPE IS VARBINARY (L)
```

for a given <Fortran host identifier> *HVN* shall be replaced by

```
INTEGER HVN_RESERVED
INTEGER HVN_LENGTH
CHARACTER HVN_DATA (* L);
```

in any <Fortran VARBINARY variable>.

h) The syntax

```
SQL TYPE IS UDTN AS PDT
```

shall be replaced by

```
ADT
```

in any <Fortran user-defined type variable>, where *ADT* is the data type listed in the “Fortran data type” column corresponding to the row for SQL data type *PDT* in Table 19, “Data type correspondences for Fortran”. *ADT* shall not be “none”. The data type identified by *UDTN* is called the *associated user-defined type* of the host variable and the data type identified by *PDT* is called the *associated SQL data type* of the host variable.

i) The syntax

```
SQL TYPE IS BLOB AS LOCATOR
```

shall be replaced by

```
INTEGER
```

in any <Fortran BLOB locator variable>. The host variable defined by <Fortran BLOB locator variable> is called a *binary large object locator variable*.

j) The syntax

```
SQL TYPE IS CLOB AS LOCATOR
```

shall be replaced by

```
INTEGER
```

in any <Fortran CLOB locator variable>. The host variable defined by <Fortran CLOB locator variable> is called a *character large object locator variable*.

k) The syntax

```
SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR
```

shall be replaced by

```
INTEGER
```

in any <Fortran user-defined type locator variable>. The host variable defined by <Fortran user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

l) The syntax

## 21.6 &lt;embedded SQL Fortran program&gt;

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

INTEGER

in any <Fortran array locator variable>. The host variable defined by <Fortran array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

## m) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

INTEGER

in any <Fortran multiset locator variable>. The host variable defined by <Fortran multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

n) 09 The syntax

SQL TYPE IS <reference type>

for a given <Fortran host identifier> *HVN* shall be replaced by

CHARACTER *HVN* \* *RL*

in any <Fortran REF variable>, where the value of *RL* is the length in octets of the reference type.

The modified <Fortran variable definition> shall be a valid Fortran type-statement in the program derived from the <embedded SQL Fortran program>.

- 7) The reference type identified by <reference type> contained in an <Fortran REF variable> is called the *referenced type* of the reference.
- 8) CHARACTER without “KIND=*n*” describes a character string variable whose equivalent SQL data type is CHARACTER with the same length and character set specified by <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit.
- 9) 14 CHARACTER KIND=*n* describes a character string variable whose equivalent SQL data type is either CHARACTER or NATIONAL CHARACTER with the same length and character set specified by <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit. The value of *n* determines implementation-defined characteristics of the Fortran variable; values of *n* are implementation-defined.
- 10) INTEGER describes an exact numeric variable. The equivalent SQL data type is INTEGER.
- 11) REAL describes an approximate numeric variable. The equivalent SQL data type is REAL.

- 12) DOUBLE PRECISION describes an approximate numeric variable. The equivalent SQL data type is DOUBLE PRECISION.
- 13) LOGICAL describes a boolean variable. The equivalent SQL data type is BOOLEAN.

## Access Rules

*None.*

## General Rules

- 1) See Subclause 21.1, “<embedded SQL host program>”.

## Conformance Rules

- 1) Without Feature B014, “Embedded Fortran”, conforming SQL language shall not contain an <embedded SQL Fortran program>.
- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <Fortran REF variable>.
- 3) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <Fortran user-defined type variable>.
- 4) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <Fortran array locator variable>.
- 5) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <Fortran multiset locator variable>.
- 6) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <Fortran user-defined type locator variable> that identifies a structured type.
- 7) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran BLOB variable>.
- 8) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran CLOB variable>.
- 9) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran BLOB locator variable>.
- 10) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran CLOB locator variable>.
- 11) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <Fortran BINARY variable>.
- 12) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <Fortran VARBINARY variable>.

## 21.7 <embedded SQL MUMPS program>

*This Subclause is modified by Subclause 18.5, “<embedded SQL MUMPS program>”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 18.6, “<embedded SQL MUMPS program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL MUMPS program>.

### Format

```

<embedded SQL MUMPS program> ::=
 !! See the Syntax Rules.

<MUMPS variable definition> ::=
 <MUMPS numeric variable> <semicolon>
 | <MUMPS character variable> <semicolon>
 | <MUMPS derived type specification> <semicolon>

<MUMPS character variable> ::=
 VARCHAR <MUMPS host identifier> <MUMPS length specification>
 [{ <comma> <MUMPS host identifier> <MUMPS length specification> }...]

<MUMPS host identifier> ::=
 !! See the Syntax Rules.

<MUMPS length specification> ::=
 <left paren> <character length> <right paren>

<MUMPS numeric variable> ::=
 <MUMPS type specification> <MUMPS host identifier>
 [{ <comma> <MUMPS host identifier> }...]

<MUMPS type specification> ::=
 INT
 | DEC [<left paren> <precision> [<comma> <scale>] <right paren>]
 | REAL

09 14 <MUMPS derived type specification> ::=
 <MUMPS CLOB variable>
 | <MUMPS VARBINARY variable>
 | <MUMPS BLOB variable>
 | <MUMPS user-defined type variable>
 | <MUMPS CLOB locator variable>
 | <MUMPS BLOB locator variable>
 | <MUMPS user-defined type locator variable>
 | <MUMPS array locator variable>
 | <MUMPS multiset locator variable>
 | <MUMPS REF variable>

<MUMPS CLOB variable> ::=
 SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<MUMPS VARBINARY variable> ::=

```

```
SQL TYPE IS VARBINARY <left paren> <length> <right paren>

<MUMPS BLOB variable> ::=
 SQL TYPE IS BLOB <left paren> <large object length> <right paren>

<MUMPS user-defined type variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS <predefined type>

<MUMPS CLOB locator variable> ::=
 SQL TYPE IS CLOB AS LOCATOR

<MUMPS BLOB locator variable> ::=
 SQL TYPE IS BLOB AS LOCATOR

<MUMPS user-defined type locator variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

<MUMPS array locator variable> ::=
 SQL TYPE IS <array type> AS LOCATOR

<MUMPS multiset locator variable> ::=
 SQL TYPE IS <multiset type> AS LOCATOR

<MUMPS REF variable> ::=
 SQL TYPE IS <reference type>
```

## Syntax Rules

- 1) An <embedded SQL MUMPS program> is a compilation unit that consists of M text and SQL text. The M text shall conform to [ISO11756]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.
- 2) A <MUMPS host identifier> is any valid M variable name. A <MUMPS host identifier> shall be contained in an <embedded SQL MUMPS program>.
- 3) An <embedded SQL statement> may be specified wherever an M command may be specified.
- 4) A <MUMPS variable definition> defines one or more host variables.
- 5) The <MUMPS character variable> describes a variable-length character string. The equivalent SQL data type is CHARACTER VARYING whose maximum length is the <character length> of the <MUMPS length specification> and whose character set is implementation-defined.
- 6) INT describes an exact numeric variable. The equivalent SQL data type is INTEGER.
- 7) DEC describes an exact numeric variable. The <scale> shall not be greater than the <precision>. The equivalent SQL data type is DECIMAL with the same <precision> and <scale>.
- 8) REAL describes an approximate numeric variable. The equivalent SQL data type is REAL.
- 9) A <MUMPS derived type specification> shall be modified as follows before it is placed into the program derived from the <embedded SQL MUMPS program> (see the Syntax Rules of [Subclause 21.1](#), “<embedded SQL host program>”).
  - a) Any optional CHARACTER SET specification shall be removed from a <MUMPS CLOB variable>.
  - b) 14 The syntax

## 21.7 &lt;embedded SQL MUMPS program&gt;

SQL TYPE IS CLOB ( *L* )

and the syntax

SQL TYPE IS BLOB ( *L* )

for a given <MUMPS host identifier> *HVN* shall be replaced by

```
INT HVN_RESERVED
INT HVN_LENGTH
VARCHAR HVN_DATA NV
```

in any <MUMPS CLOB variable> or <MUMPS BLOB variable>, where the value of *NV* is the numeric value of *L* as specified in [Subclause 6.1](#), “<data type>”.

## c) The syntax

SQL TYPE IS VARBINARY ( *L* )

for a given <MUMPS host identifier> *HVN* shall be replaced by

```
INT HVN_RESERVED
INT HVN_LENGTH
VARCHAR HVN_DATA L
```

in any <MUMPS VARBINARY variable>, where *L* is the numeric value of <length> as specified in [Subclause 5.2](#), “<token> and <separator>”.

## d) The syntax

SQL TYPE IS UDTN AS PDT

shall be replaced by

*ADT*

in any <MUMPS user-defined type variable>, where *ADT* is the data type listed in the “M data type” column corresponding to the row for SQL data type *PDT* in [Table 20](#), “[Data type correspondences for M](#)”, *ADT* shall not be “none”. The data type identified by *UDTN* is called the *associated user-defined type* of the host variable and the data type identified by *PDT* is called the *associated SQL data type* of the host variable.

## e) The syntax

SQL TYPE IS BLOB AS LOCATOR

shall be replaced by

INT

in any <MUMPS BLOB locator variable>. The host variable defined by <MUMPS BLOB locator variable> is called a *binary large object locator variable*.

## f) The syntax



SQL TYPE IS CLOB AS LOCATOR

shall be replaced by

INT

in any <MUMPS CLOB locator variable>. The host variable defined by <MUMPS CLOB locator variable> is called a *character large object locator variable*.

g) The syntax

SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

shall be replaced by

INT

in any <MUMPS user-defined type locator variable>. The host variable defined by <MUMPS user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

h) The syntax

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

INT

in any <MUMPS array locator variable>. The host variable defined by <MUMPS array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

i) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

INT

in any <MUMPS multiset locator variable>. The host variable defined by <MUMPS multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

j) 09 The syntax

SQL TYPE IS <reference type>

for a given <MUMPS host identifier> *HVN* shall be replaced by

VARCHAR *HVN* *L*

in any <MUMPS REF variable>, where the value of *L* is the length in octets of the reference type.

The modified <MUMPS variable definition> shall be a valid M variable in the program derived from the <embedded SQL MUMPS program>.

- 10) The reference type identified by <reference type> contained in an <MUMPS REF variable> is called the *referenced type* of the reference.

## Access Rules

*None.*

## General Rules

- 1) See Subclause 21.1, “<embedded SQL host program>”.

## Conformance Rules

- 1) Without Feature B015, “Embedded MUMPS”, conforming SQL language shall not contain an <embedded SQL MUMPS program>.
- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <MUMPS REF variable>.
- 3) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <MUMPS user-defined type variable>.
- 4) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <MUMPS array locator variable>.
- 5) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <MUMPS multiset locator variable>.
- 6) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <MUMPS user-defined type locator variable> that identifies a structured type.
- 7) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS BLOB variable>.
- 8) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS CLOB variable>.
- 9) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS BLOB locator variable>.
- 10) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS CLOB locator variable>.
- 11) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <MUMPS VARBINARY variable>.

## 21.8 <embedded SQL Pascal program>

*This Subclause is modified by Subclause 18.6, “<embedded SQL Pascal program>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 18.7, “<embedded SQL Pascal program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL Pascal program>.

### Format

```
<embedded SQL Pascal program> ::=
 !! See the Syntax Rules.

<Pascal variable definition> ::=
 <Pascal host identifier> [{ <comma> <Pascal host identifier> }...] <colon>
 <Pascal type specification> <semicolon>

<Pascal host identifier> ::=
 !! See the Syntax Rules.

<Pascal type specification> ::=
 PACKED ARRAY <left bracket> 1 <double period> <character length> <right bracket>
 OF CHAR [CHARACTER SET [IS] <character set specification>]
 | INTEGER
 | REAL
 | CHAR [CHARACTER SET [IS] <character set specification>]
 | BOOLEAN
 | <Pascal derived type specification>

09 14 <Pascal derived type specification> ::=
 <Pascal CLOB variable>
 | <Pascal BINARY variable>
 | <Pascal BLOB variable>
 | <Pascal user-defined type variable>
 | <Pascal CLOB locator variable>
 | <Pascal BLOB locator variable>
 | <Pascal user-defined type locator variable>
 | <Pascal array locator variable>
 | <Pascal multiset locator variable>
 | <Pascal REF variable>

<Pascal CLOB variable> ::=
 SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<Pascal BINARY variable> ::=
 SQL TYPE IS BINARY <left paren> <length> <right paren>

<Pascal BLOB variable> ::=
 SQL TYPE IS BLOB <left paren> <large object length> <right paren>

<Pascal CLOB locator variable> ::=
 SQL TYPE IS CLOB AS LOCATOR
```

**21.8 <embedded SQL Pascal program>**

```

<Pascal user-defined type variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS <predefined type>

<Pascal BLOB locator variable> ::=
 SQL TYPE IS BLOB AS LOCATOR

<Pascal user-defined type locator variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

<Pascal array locator variable> ::=
 SQL TYPE IS <array type> AS LOCATOR

<Pascal multiset locator variable> ::=
 SQL TYPE IS <multiset type> AS LOCATOR

<Pascal REF variable> ::=
 SQL TYPE IS <reference type>

```

**Syntax Rules**

- 1) An <embedded SQL Pascal program> is a compilation unit that consists of Pascal text and SQL text. The Pascal text shall conform to one of [ISO7185] or [ISO10206]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.
- 2) An <embedded SQL statement> may be specified wherever a Pascal statement may be specified. An <embedded SQL statement> may be prefixed by a Pascal label.
- 3) A <Pascal host identifier> is a Pascal variable-identifier whose applied instance denotes a defining instance within an <embedded SQL begin declare> and an <embedded SQL end declare>.
- 4) A <Pascal variable definition> defines one or more <Pascal host identifier>s.
- 5) A <Pascal variable definition> shall be modified as follows before it is placed into the program derived from the <embedded SQL Pascal program> (see the Syntax Rules of [Subclause 21.1](#), “<embedded SQL host program>”).
  - a) 14 Any optional CHARACTER SET specification shall be removed from the PACKED ARRAY OF CHAR or CHAR alternatives of a <Pascal type specification> and a <Pascal CLOB variable>.

Let  $k$  be

Case:

- i) If OCTETS is specified, then 1 (one).
  - ii) If a <character set specification> is specified, then the maximum number of octets per character of the character set identified by the <character set specification>.
  - iii) Otherwise, the maximum number of octets per character of the implementation-defined default character set.
- b) Any optional <char length units> shall be removed from a <character length> or a <character large object length>.
  - c) The <character length>  $CL$  specified in the PACKED ARRAY OF CHAR alternative of any <Pascal type specification> shall be replaced by a length equal to  $NV \cdot k$ , where  $NV$  is the numeric value of the <length> contained in  $CL$ .

- d) If any <Pascal type specification> specifies the syntax “CHAR” and if  $k$  is greater than 1 (one), then “CHAR” shall be replaced by “PACKED ARRAY [1.. $k$ ] OF CHAR”.

- e) The syntax

```
SQL TYPE IS CLOB (L)
```

for a given <Pascal host identifier>  $HVN$  shall be replaced by

```
VAR HVN = RECORD
 $HVN_RESERVED$: INTEGER;
 HVN_LENGTH : INTEGER;
 HVN_DATA : PACKED ARRAY [1.. LL] OF CHAR;
END;
```

in any <Pascal CLOB variable>, where the value of  $LL$  is  $NV*k$  and  $NV$  is the numeric value of  $L$  as defined in Subclause 6.1, “<data type>”.

- f) 14 The syntax

```
SQL TYPE IS BLOB (L)
```

for a given <Pascal host identifier>  $HVN$  shall be replaced by

```
VAR HVN = RECORD
 $HVN_RESERVED$: INTEGER;
 HVN_LENGTH : INTEGER;
 HVN_DATA : PACKED ARRAY [1.. L] OF CHAR;
END;
```

in any <Pascal BLOB variable>.

- g) The syntax

```
SQL TYPE IS BINARY (L)
```

for a given <Pascal host identifier>  $HVN$  shall be replaced by

```
 HVN : PACKED ARRAY [1.. L] OF CHAR
```

in any <Pascal BINARY variable>, where  $L$  is the numeric value of <length> as specified in Subclause 5.2, “<token> and <separator>”.

- h) The syntax

```
SQL TYPE IS $UDTN$ AS PDT
```

shall be replaced by

```
 ADT
```

in any <Pascal user-defined type variable>, where  $ADT$  is the data type listed in the “Pascal data type” column corresponding to the row for SQL data type  $PDT$  in Table 21, “Data type correspondences for Pascal”.  $ADT$  shall not be “none”. The data type identified by  $UDTN$  is called the *associated user-defined type* of the host variable and the data type identified by  $PDT$  is called the *associated SQL data type* of the host variable.

21.8 <embedded SQL Pascal program>

- i) The syntax

SQL TYPE IS BLOB AS LOCATOR

shall be replaced by

INTEGER

in any <Pascal BLOB locator variable>. The host variable defined by <Pascal BLOB locator variable> is called a *binary large object locator variable*.

- j) The syntax

SQL TYPE IS CLOB AS LOCATOR

shall be replaced by

INTEGER

in any <Pascal CLOB locator variable>. The host variable defined by <Pascal CLOB locator variable> is called a *character large object locator variable*.

- k) The syntax

SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

shall be replaced by

INTEGER

in any <Pascal user-defined type locator variable>. The host variable defined by <Pascal user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

- l) The syntax

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

INTEGER

in any <Pascal array locator variable>. The host variable defined by <Pascal array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

- m) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

INTEGER

in any <Pascal multiset locator variable>. The host variable defined by <Pascal multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

n) 09 The syntax

SQL TYPE IS <reference type>

for a given <Pascal host identifier> *HVN* shall be replaced by

*HVN* : PACKED ARRAY [1..*RL*] OF CHAR

in any <Pascal REF variable>, where the value of *RL* is the length in octets of the reference type.

The modified <Pascal variable definition> shall be a valid Pascal variable-declaration in the program derived from the <embedded SQL Pascal program>.

- 6) The reference type identified by <reference type> contained in an <Pascal REF variable> is called the *referenced type* of the reference.
- 7) CHAR specified without a CHARACTER SET specification is the ordinal-type-identifier of PASCAL. The equivalent SQL data type is CHARACTER with length 1 (one).
- 8) PACKED ARRAY [1.<character length>] OF CHAR describes a character string having two or more units (either characters or octets, as specified by <char length units>) of the simple type CHAR. The equivalent SQL data type is CHARACTER with the same length and character set specified by <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit.
- 9) INTEGER describes an exact numeric variable. The equivalent SQL data type is INTEGER.
- 10) REAL describes an approximate numeric variable. The equivalent SQL data type is REAL.
- 11) BOOLEAN describes a boolean variable. The equivalent SQL data type is BOOLEAN.

## Access Rules

*None.*

## General Rules

- 1) See Subclause 21.1, “<embedded SQL host program>”.

## Conformance Rules

- 1) Without Feature B016, “Embedded Pascal”, conforming SQL language shall not contain an <embedded SQL Pascal program>.
- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <Pascal REF variable>.
- 3) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <Pascal user-defined type variable>.
- 4) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <Pascal array locator variable>.

- 5) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <Pascal multiset locator variable>.
- 6) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <Pascal user-defined type locator variable> that identifies a structured type.
- 7) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal BLOB variable>.
- 8) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal CLOB variable>.
- 9) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal BLOB locator variable>.
- 10) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal BLOB variable>, <Pascal CLOB variable>, <Pascal CLOB locator variable>.
- 11) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <Pascal BINARY variable>.



## 21.9 <embedded SQL PL/I program>

*This Subclause is modified by Subclause 18.7, “<embedded SQL PL/I program>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 18.8, “<embedded SQL PL/I program>”, in ISO/IEC 9075-14.*

### Function

Specify an <embedded SQL PL/I program>.

### Format

```
<embedded SQL PL/I program> ::=
 !! See the Syntax Rules.

<PL/I variable definition> ::=
 { DCL | DECLARE } <PL/I type specification> [<character representation>...] <semicolon>
 | { <PL/I host identifier> | <left paren> <PL/I host identifier>
 [{ <comma> <PL/I host identifier> }...] <right paren> }
 <PL/I type specification> [<character representation>...] <semicolon>

<PL/I host identifier> ::=
 !! See the Syntax Rules.

<PL/I type specification> ::=
 { CHAR | CHARACTER } [VARYING] <left paren> <character length> <right paren>
 [CHARACTER SET [IS] <character set specification>]
 | <PL/I type fixed decimal> <left paren> <precision> [<comma> <scale>] <right paren>
 | <PL/I type fixed binary> [<left paren> <precision> <right paren>]
 | <PL/I type float binary> <left paren> <precision> <right paren>
 | <PL/I derived type specification>

09 14 <PL/I derived type specification> ::=
 <PL/I CLOB variable>
 | <PL/I BINARY variable>
 | <PL/I VARBINARY variable>
 | <PL/I BLOB variable>
 | <PL/I user-defined type variable>
 | <PL/I CLOB locator variable>
 | <PL/I BLOB locator variable>
 | <PL/I user-defined type locator variable>
 | <PL/I array locator variable>
 | <PL/I multiset locator variable>
 | <PL/I REF variable>

<PL/I CLOB variable> ::=
 SQL TYPE IS CLOB <left paren> <character large object length> <right paren>
 [CHARACTER SET [IS] <character set specification>]

<PL/I BINARY variable> ::=
 SQL TYPE IS BINARY <left paren> <length> <right paren>

<PL/I VARBINARY variable> ::=
 SQL TYPE IS VARBINARY <left paren> <length> <right paren>

<PL/I BLOB variable> ::=
```

```

 SQL TYPE IS BLOB <left paren> <large object length> <right paren>

<PL/I user-defined type variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS <predefined type>

<PL/I CLOB locator variable> ::=
 SQL TYPE IS CLOB AS LOCATOR

<PL/I BLOB locator variable> ::=
 SQL TYPE IS BLOB AS LOCATOR

<PL/I user-defined type locator variable> ::=
 SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

<PL/I array locator variable> ::=
 SQL TYPE IS <array type> AS LOCATOR

<PL/I multiset locator variable> ::=
 SQL TYPE IS <multiset type> AS LOCATOR

<PL/I REF variable> ::=
 SQL TYPE IS <reference type>

<PL/I type fixed decimal> ::=
 { DEC | DECIMAL } FIXED
 | FIXED { DEC | DECIMAL }

<PL/I type fixed binary> ::=
 { BIN | BINARY } FIXED
 | FIXED { BIN | BINARY }

<PL/I type float binary> ::=
 { BIN | BINARY } FLOAT
 | FLOAT { BIN | BINARY }

```

## Syntax Rules

- 1) An <embedded SQL PL/I program> is a compilation unit that consists of PL/I text and SQL text. The PL/I text shall conform to [ISO6160]. The SQL text shall consist of one or more <embedded SQL statement>s and, optionally, one or more <embedded SQL declare section>s.
- 2) An <embedded SQL statement> may be specified wherever a PL/I statement may be specified within a procedure block. If the PL/I statement could include a label prefix, the <embedded SQL statement> may be immediately preceded by a label prefix.
- 3) A <PL/I host identifier> is any valid PL/I variable identifier. A <PL/I host identifier> shall be contained in an <embedded SQL PL/I program>.
- 4) A <PL/I variable definition> defines one or more host variables.
- 5) A <PL/I variable definition> shall be modified as follows before it is placed into the program derived from the <embedded SQL PL/I program> (see the Syntax Rules of [Subclause 21.1](#), “<embedded SQL host program>”).
  - a) Any optional CHARACTER SET specification shall be removed from the CHARACTER or CHARACTER VARYING alternatives of a <PL/I type specification>.

Let  $k$  be

Case:

- i) If OCTETS is specified, then 1 (one).
  - ii) If a <character set specification> is specified, then the maximum number of octets per character of the character set identified by the <character set specification>.
  - iii) Otherwise, the maximum number of octets per character of the implementation-defined default character set.
- b) Any optional <char length units> shall be removed from a <character length> or a <character large object length>.
- c) The <character length> *CL* specified in the CHARACTER or CHARACTER VARYING alternatives of any <PL/I type specification> shall be replaced by a length equal to  $NV*k$ , where *NV* is the numeric value of the <length> contained in *CL*.
- d) 14 The syntax

SQL TYPE IS CLOB ( *L* )

for a given <PL/I host identifier> *HVN* shall be replaced by

```
DCL 1 HVN
 2 HVN_RESERVED FIXED BINARY(31),
 2 HVN_LENGTH FIXED BINARY(31),
 2 HVN_DATA CHARACTER(LL);
```

in any <PL/I CLOB variable>, where the value of *LL* is  $NV*k$  and *NV* is the numeric value of *L* as defined in Subclause 6.1, “<data type>”.

- e) The syntax

SQL TYPE IS BLOB ( *L* )

for a given <PL/I host identifier> *HVN* shall be replaced by

```
DCL 1 HVN
 2 HVN_RESERVED FIXED BINARY(31),
 2 HVN_LENGTH FIXED BINARY(31),
 2 HVN_DATA CHARACTER(L);
```

in any <PL/I BLOB variable>.

- f) The syntax

SQL TYPE IS BINARY ( *L* )

for a given <PL/I host identifier> *HVN* shall be replaced by

```
DCL HVN CHARACTER(L)
```

in any <PL/I BINARY variable>.

- g) The syntax

SQL TYPE IS VARBINARY ( L )

for a given <PL/I host identifier> *HVN* shall be replaced by

DCL *HVN* CHARACTER(L) VARYING

in any <PL/I VARBINARY variable>.

h) The syntax

SQL TYPE IS *UDTN* AS *PDT*

shall be replaced by

*ADT*

in any <PL/I user-defined type variable>, where *ADT* is the data type listed in the “PL/I data type” column corresponding to the row for SQL data type *PDT* in Table 22, “Data type correspondences for PL/I”. *ADT* shall not be “none”. The data type identified by *UDTN* is called the *associated user-defined type* of the host variable and the data type identified by *PDT* is called the *associated SQL data type* of the host variable.

i) The syntax

SQL TYPE IS BLOB AS LOCATOR

shall be replaced by

FIXED BINARY(31)

in any <PL/I BLOB locator variable>. The host variable defined by <PL/I BLOB locator variable> is called a *binary large object locator variable*.

j) The syntax

SQL TYPE IS CLOB AS LOCATOR

shall be replaced by

FIXED BINARY(31)

in any <PL/I CLOB locator variable>. The host variable defined by <PL/I CLOB locator variable> is called a *character large object locator variable*.

k) The syntax

SQL TYPE IS <path-resolved user-defined type name> AS LOCATOR

shall be replaced by

FIXED BINARY(31)

in any <PL/I user-defined type locator variable>. The host variable defined by <PL/I user-defined type locator variable> is called a *user-defined type locator variable*. The data type identified by <path-resolved user-defined type name> is called the *associated user-defined type* of the host variable.

l) The syntax

SQL TYPE IS <array type> AS LOCATOR

shall be replaced by

FIXED BINARY(31)

in any <PL/I array locator variable>. The host variable defined by <PL/I array locator variable> is called an *array locator variable*. The data type identified by <array type> is called the *associated array type* of the host variable.

m) The syntax

SQL TYPE IS <multiset type> AS LOCATOR

shall be replaced by

FIXED BINARY(31)

in any <PL/I multiset locator variable>. The host variable defined by <PL/I multiset locator variable> is called a *multiset locator variable*. The data type identified by <multiset type> is called the *associated multiset type* of the host variable.

n) 09 The syntax

SQL TYPE IS <reference type>

for a given <PL/I host identifier> *HVN* shall be replaced by

DCL *HVN* CHARACTER(*RL*) VARYING

in any <PL/I REF variable>, where the value of *RL* is the length in octets of the reference type.

The modified <PL/I variable definition> shall be a valid PL/I data declaration in the program derived from the <embedded SQL PL/I program>.

- 6) The reference type identified by <reference type> contained in an <PL/I REF variable> is called the *referenced type* of the reference.
- 7) A <PL/I variable definition> shall specify a scalar variable, not an array or structure.
- 8) The optional <character representation> sequence in a <PL/I variable definition> may specify an INITIAL clause. Whether other clauses may be specified is implementation-defined. The <character representation> sequence shall be such that the <PL/I variable definition> is a valid PL/I DECLARE statement.
- 9) 14 CHARACTER describes a character string variable whose equivalent SQL data type has the character set specified by <character set specification>. If <character set specification> is not specified, then an implementation-defined <character set specification> is implicit.

Case:

- a) If VARYING is not specified, then the length of the variable is fixed. The equivalent SQL data type is CHARACTER with the same length, measured in the units specified by <char length units>.
- b) If VARYING is specified, then the variable is of variable length, with maximum size the value of <length>. The equivalent SQL data type is CHARACTER VARYING with the same maximum length, measured in the units specified by <char length units>.

- 10) 14 FIXED DECIMAL describes an exact numeric variable. The <scale> shall not be greater than the <precision>. The equivalent SQL data type is DECIMAL with the same <precision> and <scale>.
- 11) FIXED BINARY describes an exact numeric variable. The equivalent SQL data type is SMALLINT, INTEGER, or BIGINT.
- 12) FLOAT BINARY describes an approximate numeric variable. The equivalent SQL data type is FLOAT with the same <precision>.

## Access Rules

*None.*

## General Rules

- 1) See Subclause 21.1, “<embedded SQL host program>”

## Conformance Rules

- 1) Without Feature B017, “Embedded PL/I”, conforming SQL language shall not contain an <embedded SQL PL/I program>.
- 2) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <PL/I REF variable>.
- 3) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <PL/I user-defined type variable>.
- 4) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <PL/I array locator variable>.
- 5) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <PL/I multiset locator variable>.
- 6) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <PL/I user-defined type locator variable> that identifies a structured type.
- 7) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I BLOB variable>.
- 8) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I CLOB variable>.
- 9) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I BLOB locator variable>.
- 10) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I CLOB locator variable>.
- 11) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <PL/I BINARY variable>.

- 12) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <PL/I VARBINARY variable>.

*(Blank page)*



## 22 Direct invocation of SQL

### 22.1 <direct SQL statement>

#### Function

Specify direct execution of SQL.

#### Format

```
<direct SQL statement> ::=
 <directly executable statement> <semicolon>

<directly executable statement> ::=
 <direct SQL data statement>
 | <SQL schema statement>
 | <SQL transaction statement>
 | <SQL connection statement>
 | <SQL session statement>
 | <direct implementation-defined statement>

<direct SQL data statement> ::=
 <delete statement: searched>
 | <direct select statement: multiple rows>
 | <insert statement>
 | <update statement: searched>
 | <truncate table statement>
 | <merge statement>
 | <temporary table declaration>

<direct implementation-defined statement> ::=
 !! See the Syntax Rules.
```

#### Syntax Rules

- 1) The <direct SQL data statement> shall not contain an SQL parameter reference, SQL variable reference, <dynamic parameter specification>, or <embedded variable specification>.
- 2) The <value specification> that represents the null value is implementation-defined.
- 3) The Format and Syntax Rules for <direct implementation-defined statement> are implementation-defined.

#### Access Rules

- 1) The Access Rules for <direct implementation-defined statement> are implementation-defined.

## General Rules

- 1) The following <direct SQL statement>s are transaction-initiating <direct SQL statement>s:
  - a) <direct SQL statement>s that are transaction-initiating <SQL procedure statement>s.
  - b) <direct select statement: multiple rows>.
  - c) <direct implementation-defined statement>s that are transaction-initiating.
- 2) After the last invocation of an SQL-statement by an SQL-agent in an SQL-session:
  - a) A <rollback statement> or a <commit statement> is effectively executed. If an unrecoverable error has occurred, or if the direct invocation of SQL terminated unexpectedly, or if any constraint is not satisfied, then a <rollback statement> is performed. Otherwise, the choice of which of these SQL-statements to perform is implementation-dependent. The determination of whether a direct invocation of SQL has terminated unexpectedly is implementation-dependent.
  - b) Let *D* be the <descriptor name> of any SQL descriptor area that is currently allocated within the current SQL-session. A <deallocate descriptor statement> that specifies  
  
DEALLOCATE DESCRIPTOR *D*  
  
is effectively executed.
  - c) All SQL-sessions associated with the SQL-agent are terminated.
- 3) A copy of the top cell of the authorization stack is pushed onto the authorization stack.
- 4) Let *S* be the <direct SQL statement>.
- 5) If *S* does not conform to the Format, Syntax Rules, and Access Rules for a <direct SQL statement>, then an exception condition is raised: *syntax error or access rule violation*.
- 6) When *S* is invoked by the SQL-agent,  
Case:
  - a) If *S* is an <SQL connection statement>, then:
    - i) The first diagnostics area is emptied.
    - ii) *S* is executed.
    - iii) If *S* successfully initiated or resumed an SQL-session, then subsequent invocations of a <direct SQL statement> by the SQL-agent are associated with that SQL-session until the SQL-agent terminates the SQL-session or makes it dormant.
  - b) Otherwise:
    - i) If no SQL-session is current for the SQL-agent, then  
Case:
      - 1) If the SQL-agent has not executed an <SQL connection statement> and there is no default SQL-session associated with the SQL-agent, then the following <connect statement> is effectively executed:

CONNECT TO DEFAULT

- 2) If the SQL-agent has not executed an <SQL connection statement> and there is a default SQL-session associated with the SQL-agent, then the following <set connection statement> is effectively executed:

SET CONNECTION DEFAULT

- 3) Otherwise, an exception condition is raised: *connection exception — connection does not exist*.

Subsequent calls to an <externally-invoked procedure> or invocations of a <direct SQL statement> by the SQL-agent are associated with the SQL-session until the SQL-agent terminates the SQL-session or makes it dormant.

- ii) If an SQL-transaction is active for the SQL-agent, then *S* is associated with that SQL-transaction. If *S* is a <direct implementation-defined statement>, then it is implementation-defined whether or not *S* may be associated with an active SQL-transaction; if not, then an exception condition is raised: *invalid transaction state — active SQL-transaction*.
- iii) If no SQL-transaction is active for the SQL-agent, then
  - 1) Case:
    - A) If *S* is a transaction-initiating <direct SQL statement>, then an SQL-transaction is initiated.
    - B) If *S* is a <direct implementation-defined statement>, then it is implementation-defined whether or not *S* initiates an SQL-transaction. If an implementation defines *S* to be transaction-initiating, then an SQL-transaction is initiated.
  - 2) If *S* initiated an SQL-transaction, then:
    - A) Let *T* be the SQL-transaction initiated by *S*.
    - B) *T* is associated with this invocation and any subsequent invocations of <direct SQL statement>s or calls to an <externally-invoked procedure> by the SQL-agent until the SQL-agent terminates *T*.
    - C) If *S* is not a <start transaction statement>, then
 

Case:

      - I) If a <set transaction statement> has been executed since the termination of the last SQL-transaction in the SQL-session (or if there has been no previous SQL-transaction in the SQL-session and a <set transaction statement> has been executed), then the access mode, constraint mode, and isolation level of *T* are set as specified by the <set transaction statement>.
      - II) Otherwise, the access mode, constraint mode for all constraints, and isolation level for *T* are *read-write*, *immediate*, and *SERIALIZABLE*, respectively.
    - D) *T* is associated with the SQL-session.
- iv) If *S* contains an <SQL schema statement> and the access mode of the current SQL-transaction is *read-only*, then an exception condition is raised: *invalid transaction state — read-only SQL-*

22.1 <direct SQL statement>

*transaction*, no further subrules of this General Rule are evaluated, and evaluation continues with the next General Rule.

v) The first diagnostics area is emptied.

vi) *S* is executed.

7) Upon completion of execution, the top cell in the authorization stack is removed.

8) If the execution of a <direct SQL data statement> occurs within the same SQL-transaction as the execution of an SQL-schema statement and this is not allowed by the SQL-implementation, then an exception condition is raised: *invalid transaction state — schema and data statement mixing not supported*.

9) Case:

a) If *S* executed successfully, then either a completion condition is raised: *successful completion*, or a completion condition is raised: *warning*, or a completion condition is raised: *no data*.

b) If *S* did not execute successfully, then all changes made to SQL-data or schemas by the execution of *S* are canceled and an exception condition is raised.

NOTE 472 — The method of raising a condition is implementation-defined.

10) Diagnostics information resulting from the execution of *S* is placed into the first diagnostics area, causing the first condition area in the first diagnostics area to become occupied.

NOTE 473 — The method of accessing the diagnostics information is implementation-defined, but does not alter the contents of the diagnostics area.

## Conformance Rules

1) Without Feature B021, “Direct SQL”, conforming SQL language shall not contain a <direct SQL statement>.

## 22.2 <direct select statement: multiple rows>

### Function

Specify a statement to retrieve multiple rows from a specified table.

### Format

```
<direct select statement: multiple rows> ::=
 <cursor specification>
```

### Syntax Rules

- 1) The <query expression> of a <direct select statement: multiple rows> shall not contain a <value specification> other than a <literal>, CURRENT\_USER, CURRENT\_ROLE, SESSION\_USER, SYSTEM\_USER, CURRENT\_CATALOG, CURRENT\_SCHEMA, CURRENT\_PATH, CURRENT\_DEFAULT\_TRANSFORM\_GROUP, or CURRENT\_TRANSFORM\_GROUP\_FOR\_TYPE.
- 2) The <cursor specification> shall not contain an <updatability clause>.

### Access Rules

*None.*

### General Rules

- 1) Let  $Q$  be the result of the <cursor specification>.
- 2) Case:
  - a) If  $Q$  is empty, then a completion condition is raised: *no data*.
  - b) Otherwise,  $Q$  is not empty and  $Q$  is returned. The method of returning  $Q$  is implementation-defined.

### Conformance Rules

*None.*

*(Blank page)*

## 23 Diagnostics management

### 23.1 <get diagnostics statement>

*This Subclause is modified by Subclause 17.1, “<get diagnostics statement>”, in ISO/IEC 9075-4.  
This Subclause is modified by Subclause 23.1, “<get diagnostics statement>”, in ISO/IEC 9075-9.  
This Subclause is modified by Subclause 19.1, “<get diagnostics statement>”, in ISO/IEC 9075-14.*

#### Function

Get exception or completion condition information from a diagnostics area.

#### Format

```
04 <get diagnostics statement> ::=
 GET DIAGNOSTICS <SQL diagnostics information>

<SQL diagnostics information> ::=
 <statement information>
 | <condition information>
 | <all information>

<statement information> ::=
 <statement information item> [{ <comma> <statement information item> }...]

<statement information item> ::=
 <simple target specification> <equals operator> <statement information item name>

<statement information item name> ::=
 NUMBER
 | MORE
 | COMMAND_FUNCTION
 | COMMAND_FUNCTION_CODE
 | DYNAMIC_FUNCTION
 | DYNAMIC_FUNCTION_CODE
 | ROW_COUNT
 | TRANSACTIONS_COMMITTED
 | TRANSACTIONS_ROLLED_BACK
 | TRANSACTION_ACTIVE

<condition information> ::=
 CONDITION <condition number> <condition information item>
 [{ <comma> <condition information item> }...]

<condition information item> ::=
 <simple target specification> <equals operator> <condition information item name>

04 <condition information item name> ::=
```

## 23.1 &lt;get diagnostics statement&gt;

```

 CATALOG_NAME
 | CLASS_ORIGIN
 | COLUMN_NAME
 | CONDITION_NUMBER
 | CONNECTION_NAME
 | CONSTRAINT_CATALOG
 | CONSTRAINT_NAME
 | CONSTRAINT_SCHEMA
 | CURSOR_NAME
 | MESSAGE_LENGTH
 | MESSAGE_OCTET_LENGTH
 | MESSAGE_TEXT
 | PARAMETER_MODE
 | PARAMETER_NAME
 | PARAMETER_ORDINAL_POSITION
 | RETURNED_SQLSTATE
 | ROUTINE_CATALOG
 | ROUTINE_NAME
 | ROUTINE_SCHEMA
 | SCHEMA_NAME
 | SERVER_NAME
 | SPECIFIC_NAME
 | SUBCLASS_ORIGIN
 | TABLE_NAME
 | TRIGGER_CATALOG
 | TRIGGER_NAME
 | TRIGGER_SCHEMA

```

```

<all information> ::=
 <all info target> <equals operator> ALL [<all qualifier>]

<all info target> ::=
 <simple target specification>

<all qualifier> ::=
 STATEMENT
 | CONDITION [<condition number>]

<condition number> ::=
 <simple value specification>

```

## Syntax Rules

- 1) Let the <simple target specification> contained in a <statement information item> or <condition information item> be *STS*. Let *IDDT* be the declared type specified in Table 30, “<statement information item name>s for use with <get diagnostics statement>”, for the corresponding <statement information item name> or in Table 31, “<condition information item name>s for use with <get diagnostics statement>”, for the corresponding <condition information item name>. Let *ID* be the expression:

```
CAST (' ' AS IDDT)
```

*ID* shall be assignable to *STS* according to the Syntax Rules of Subclause 9.1, “Retrieval assignment”, with *ID* as *VALUE* and *STS* as *TARGET*.

- 2) The declared type of <all info target> shall be a character string type.



- 3) The declared type of <condition number> shall be exact numeric with scale 0 (zero).

**Table 30** — <statement information item name>s for use with <get diagnostics statement>

<identifier>	Declared Type
COMMAND_FUNCTION	variable-length character string with implementation-defined maximum length
COMMAND_FUNCTION_CODE	exact numeric with scale 0 (zero)
DYNAMIC_FUNCTION	variable-length character string with implementation-defined maximum length
DYNAMIC_FUNCTION_CODE	exact numeric with scale 0 (zero)
MORE	fixed-length character string with length 1 (one)
NUMBER	exact numeric with scale 0 (zero)
ROW_COUNT	exact numeric with scale 0 (zero)
TRANSACTION_ACTIVE	exact numeric with scale 0 (zero)
TRANSACTIONS_COMMITTED	exact numeric with scale 0 (zero)
TRANSACTIONS_ROLLED_BACK	exact numeric with scale 0 (zero)

**Table 31** — <condition information item name>s for use with <get diagnostics statement>

<identifier>	Declared Type
CATALOG_NAME	variable-length character string with implementation-defined maximum length
CLASS_ORIGIN	variable-length character string with implementation-defined maximum length
COLUMN_NAME	variable-length character string with implementation-defined maximum length
CONDITION_NUMBER	exact numeric with scale 0 (zero)
CONNECTION_NAME	variable-length character string with implementation-defined maximum length

<identifier>	Declared Type
CONSTRAINT_CATALOG	variable-length character string with implementation-defined maximum length
CONSTRAINT_NAME	variable-length character string with implementation-defined maximum length
CONSTRAINT_SCHEMA	variable-length character string with implementation-defined maximum length
CURSOR_NAME	variable-length character string with implementation-defined maximum length
MESSAGE_LENGTH	exact numeric with scale 0 (zero)
MESSAGE_OCTET_LENGTH	exact numeric with scale 0 (zero)
MESSAGE_TEXT	variable-length character string with implementation-defined maximum length
PARAMETER_MODE	variable-length character string with maximum length 5
PARAMETER_NAME	variable-length character string with implementation-defined maximum length
PARAMETER_ORDINAL_POSITION	exact numeric with scale 0 (zero)
RETURNED_SQLSTATE	fixed-length character string with length 5
ROUTINE_CATALOG	variable-length character string with implementation-defined maximum length
ROUTINE_NAME	variable-length character string with implementation-defined maximum length
ROUTINE_SCHEMA	variable-length character string with implementation-defined maximum length
SCHEMA_NAME	variable-length character string with implementation-defined maximum length
SERVER_NAME	variable-length character string with implementation-defined maximum length
SPECIFIC_NAME	variable-length character string with implementation-defined maximum length
SUBCLASS_ORIGIN	variable-length character string with implementation-defined maximum length

<identifier>	Declared Type
TABLE_NAME	variable-length character string with implementation-defined maximum length
TRIGGER_CATALOG	variable-length character string with implementation-defined maximum length
TRIGGER_NAME	variable-length character string with implementation-defined maximum length
TRIGGER_SCHEMA	variable-length character string with implementation-defined maximum length

## Access Rules

*None.*

## General Rules

- 1) 04 Let *DA* be the first diagnostics area.
- 2) Specification of <statement information item> assigns the value of the specified statement information item in *DA* to <simple target specification>.
  - a) The value of NUMBER is the number of exception or completion conditions that have been stored in *DA* as a result of executing the previous SQL-statement other than a <get diagnostics statement>.
 

NOTE 474 — The <get diagnostics statement> itself may return information via the SQLSTATE parameter, but does not modify the previous contents of *DA*.
  - b) The value of MORE is:

Y	More conditions were raised during execution of the SQL-statement than there are condition areas in <i>DA</i> .
N	All of the conditions that were raised during execution of the SQL-statement have been stored in <i>DA</i> .

- c) The value of COMMAND\_FUNCTION is the identification of the SQL-statement executed. Table 32, “SQL-statement codes” specifies the identifier of the SQL-statements.
- d) The value of COMMAND\_FUNCTION\_CODE is a number identifying the SQL-statement executed. Table 32, “SQL-statement codes” specifies the code for the SQL-statements. Positive values are reserved for SQL-statements defined by ISO/IEC 9075; negative values are reserved for implementation-defined SQL-statements.
- e) The value of DYNAMIC\_FUNCTION and DYNAMIC\_FUNCTION\_CODE are, respectively, a character string and a number, each of which identifies the type of the SQL-statement being prepared or executed dynamically. Table 32, “SQL-statement codes”, specifies the identifiers and codes of the

## 23.1 &lt;get diagnostics statement&gt;

SQL-statements. Positive code values are reserved for SQL-statements defined by ISO/IEC 9075; negative code values are reserved for implementation-defined SQL-statements. It is implementation-defined whether the identifier and code from Table 32, “SQL-statement codes”, for <dynamic select statement> or <dynamic single row select statement> are used to describe a <dynamic select statement> or a <dynamic single row select statement> that has been prepared but has not yet been executed dynamically.

03 04 09 14 Table 32 — SQL-statement codes

SQL-statement	Identifier	Code
<allocate cursor statement>	ALLOCATE CURSOR	1 (one)
<allocate descriptor statement>	ALLOCATE DESCRIPTOR	2
<alter domain statement>	ALTER DOMAIN	3
<alter routine statement>	ALTER ROUTINE	17
<alter sequence generator statement>	ALTER SEQUENCE	134
<alter type statement>	ALTER TYPE	60
<alter table statement>	ALTER TABLE	4
<alter transform statement>	ALTER TRANSFORM	127
<assertion definition>	CREATE ASSERTION	6
<call statement>	CALL	7
<character set definition>	CREATE CHARACTER SET	8
<close statement>	CLOSE CURSOR	9
<collation definition>	CREATE COLLATION	10
<commit statement>	COMMIT WORK	11
<connect statement>	CONNECT	13
<deallocate descriptor statement>	DEALLOCATE DESCRIPTOR	15
<deallocate prepared statement>	DEALLOCATE PREPARE	16
<delete statement: positioned>	DELETE CURSOR	18
<delete statement: searched>	DELETE WHERE	19
<describe statement>	DESCRIBE	20

SQL-statement	Identifier	Code
<direct select statement: multiple rows>	SELECT	21
<disconnect statement>	DISCONNECT	22
<domain definition>	CREATE DOMAIN	23
<drop assertion statement>	DROP ASSERTION	24
<drop character set statement>	DROP CHARACTER SET	25
<drop collation statement>	DROP COLLATION	26
<drop data type statement>	DROP TYPE	35
<drop domain statement>	DROP DOMAIN	27
<drop role statement>	DROP ROLE	29
<drop routine statement>	DROP ROUTINE	30
<drop schema statement>	DROP SCHEMA	31
<drop sequence generator statement>	DROP SEQUENCE	135
<drop table statement>	DROP TABLE	32
<drop transform statement>	DROP TRANSFORM	116
<drop transliteration statement>	DROP TRANSLATION	33
<drop trigger statement>	DROP TRIGGER	34
<drop user-defined cast statement>	DROP CAST	78
<drop user-defined ordering statement>	DROP ORDERING	115
<drop view statement>	DROP VIEW	36
<dynamic close statement>	DYNAMIC CLOSE	37
<dynamic delete statement: positioned>	DYNAMIC DELETE CURSOR	38
<dynamic fetch statement>	DYNAMIC FETCH	39
<dynamic open statement>	DYNAMIC OPEN	40
<dynamic select statement>	SELECT CURSOR	85
<dynamic single row select statement>	SELECT	41

SQL-statement	Identifier	Code
<dynamic update statement: positioned>	DYNAMIC UPDATE CURSOR	42
<execute immediate statement>	EXECUTE IMMEDIATE	43
<execute statement>	EXECUTE	44
<fetch statement>	FETCH	45
<free locator statement>	FREE LOCATOR	98
<get descriptor statement>	GET DESCRIPTOR	47
<hold locator statement>	HOLD LOCATOR	99
<grant privilege statement>	GRANT	48
<grant role statement>	GRANT ROLE	49
<insert statement>	INSERT	50
<merge statement>	MERGE	128
<open statement>	OPEN	53
<preparable dynamic delete statement: positioned>	PREPARABLE DYNAMIC DELETE CURSOR	54
<preparable dynamic update statement: positioned>	PREPARABLE DYNAMIC UPDATE CURSOR	55
<prepare statement>	PREPARE	56
<release savepoint statement>	RELEASE SAVEPOINT	57
<return statement>	RETURN	58
<revoke privilege statement>	REVOKE	59
<revoke role statement>	REVOKE ROLE	129
<role definition>	CREATE ROLE	61
<rollback statement>	ROLLBACK WORK	62
<savepoint statement>	SAVEPOINT	63
<schema definition>	CREATE SCHEMA	64
<schema routine>	CREATE ROUTINE	14

SQL-statement	Identifier	Code
<select statement: single row>	SELECT	65
<sequence generator definition>	CREATE SEQUENCE	133
<set catalog statement>	SET CATALOG	66
<set connection statement>	SET CONNECTION	67
<set constraints mode statement>	SET CONSTRAINT	68
<set descriptor statement>	SET DESCRIPTOR	70
<set local time zone statement>	SET TIME ZONE	71
<set names statement>	SET NAMES	72
<set path statement>	SET PATH	69
<set role statement>	SET ROLE	73
<set schema statement>	SET SCHEMA	74
<set session user identifier statement>	SET SESSION AUTHORIZATION	76
<set session characteristics statement>	SET SESSION CHARACTERISTICS	109
<set session collation statement>	SET COLLATION	136
<set transform group statement>	SET TRANSFORM GROUP	118
<set transaction statement>	SET TRANSACTION	75
<start transaction statement>	START TRANSACTION	111
<table definition>	CREATE TABLE	77
<transform definition>	CREATE TRANSFORM	117
<transliteration definition>	CREATE TRANSLATION	79
<trigger definition>	CREATE TRIGGER	80
<truncate table statement>	TRUNCATE TABLE	139
<update statement: positioned>	UPDATE CURSOR	81
<update statement: searched>	UPDATE WHERE	82
<user-defined cast definition>	CREATE CAST	52

SQL-statement	Identifier	Code
<user-defined type definition>	CREATE TYPE	83
<user-defined ordering definition>	CREATE ORDERING	114
<view definition>	CREATE VIEW	84
Implementation-defined statements	An implementation-defined character string value different from the value associated with any other SQL-statement	$x^1$
Unrecognized statements	A zero-length string	0 (zero)
<sup>1</sup> An implementation-defined negative number different from the value associated with any other SQL-statement.		

NOTE 475 — Other, additional, values are used in other parts of ISO/IEC 9075; please see the corresponding table in the other parts of ISO/IEC 9075; for more information.

- f) The value of ROW\_COUNT is the number of rows affected as the result of executing a <delete statement: searched>, <insert statement>, <merge statement>, or <update statement: searched> or as a direct result of executing the previous SQL-statement. Let *S* be the <delete statement: searched>, <insert statement>, <merge statement>, or <update statement: searched>. Let *T* be the table identified by the <table name> directly contained in *S*.

Case:

- i) If *S* is an <insert statement>, then the value of ROW\_COUNT is the number of rows inserted into *T*.
- ii) If *S* is a <merge statement>, then let *TR1* be the <target table> immediately contained in *S*, let *TR2* be the <table reference> immediately contained in *S*, and let *SC* be the <search condition> immediately contained in *S*. If <merge correlation name> is specified, let *MCN* be “AS <merge correlation name>”; otherwise, let *MCN* be a zero-length string.

Case:

- 1) If *S* contains a <merge when matched clause> and does not contain a <merge when not matched clause>, then the value of ROW\_COUNT is effectively derived by executing the statement:

```
SELECT COUNT (*)
FROM TR1 MCN, TR2
WHERE SC
```

before the execution of *S*.

- 2) If *S* contains a <merge when not matched clause> and does not contain a <merge when matched clause>, then the value of ROW\_COUNT is effectively derived by executing the statement:

```
(SELECT COUNT (*)
 FROM TR1 MCN
 RIGHT OUTER JOIN
```



```

 TR2
 ON SC)
-
(SELECT COUNT (*)
 FROM TR1 MCN, TR2
 WHERE SC)

```

before the execution of *S*.

- 3) If *S* contains both a <merge when matched clause> and a <merge when not matched clause>, then the value of ROW\_COUNT is effectively derived by executing the statement:

```

SELECT COUNT (*)
FROM TR1 MCN
 RIGHT OUTER JOIN
 TR2
 ON SC

```

before the execution of *S*.

- iii) If <correlation name> is specified, then let *MCN* be “AS <correlation name>”; otherwise, let *MCN* be a zero-length string. If *S* is a <delete statement: searched> or an <update statement: searched>, then

Case:

- 1) If *S* does not contain a <search condition>, then the value of ROW\_COUNT is the cardinality of *T* before the execution of *S*.
- 2) Otherwise, let *SC* be the <search condition> directly contained in *S*. The value of ROW\_COUNT is effectively derived by executing the statement:

```

SELECT COUNT (*)
FROM T MCN
WHERE SC

```

before the execution of *S*.

The value of ROW\_COUNT following the execution of an SQL-statement that does not directly result in the execution of a <delete statement: searched>, an <insert statement>, a <merge statement>, or an <update statement: searched> is implementation-dependent.

- g) The value of TRANSACTIONS\_COMMITTED is the number of SQL-transactions that have been committed since the most recent time at which *DA* was emptied.

NOTE 476 — See the General Rules of Subclause 13.3, “<externally-invoked procedure>”, and Subclause 13.4, “Calls to an <externally-invoked procedure>”. TRANSACTIONS\_COMMITTED indicates the number of SQL-transactions that were committed during the invocation of an external routine.

- h) The value of TRANSACTIONS\_ROLLED\_BACK is the number of SQL-transactions that have been rolled back since the most recent time at which *DA* was emptied.

NOTE 477 — See the General Rules of Subclause 13.3, “<externally-invoked procedure>”, and Subclause 13.4, “Calls to an <externally-invoked procedure>”. TRANSACTIONS\_ROLLED\_BACK indicates the number of SQL-transactions that were rolled back during the invocation of an external routine.

- i) The value of TRANSACTION\_ACTIVE is 1 (one) if an SQL-transaction is currently active, and 0 (zero) if an SQL-transaction is not currently active.

## 23.1 &lt;get diagnostics statement&gt;

NOTE 478 — TRANSACTION\_ACTIVE indicates whether an SQL-transaction is active upon return from an external routine.

- 3) If <condition information> is specified, then let  $N$  be the value of <condition number>. If  $N$  is less than 1 (one) or greater than the number of occupied condition areas in  $DA$ , then an exception condition is raised: *invalid condition number*. If <condition number> has the value 1 (one), then the diagnostics information retrieved corresponds to the condition indicated by the SQLSTATE value actually returned by execution of the previous SQL-statement other than a <get diagnostics statement>. Otherwise, the association between <condition number> values and specific conditions raised during evaluation of the General Rules for that SQL-statement is implementation-dependent.
- 4) Specification of <condition information item> assigns the value of the specified condition information item in the  $N$ -th condition area in  $DA$  to <simple target specification>.
  - a) The value of CONDITION\_NUMBER is the value of <condition number>.
  - b) The value of CLASS\_ORIGIN is the identification of the naming authority that defined the class value of RETURNED\_SQLSTATE. That value shall be 'ISO 9075' for any RETURNED\_SQLSTATE whose class value is fully defined in [Subclause 24.1, “SQLSTATE”](#), and shall be an implementation-defined character string other than 'ISO 9075' for any RETURNED\_SQLSTATE whose class value is an implementation-defined class value.
  - c) The value of SUBCLASS\_ORIGIN is the identification of the naming authority that defined the subclass value of RETURNED\_SQLSTATE. That value shall be 'ISO 9075' for any RETURNED\_SQLSTATE whose subclass value is fully defined in [Subclause 24.1, “SQLSTATE”](#), and shall be an implementation-defined character string other than 'ISO 9075' for any RETURNED\_SQLSTATE whose subclass value is an implementation-defined subclass value.
  - d) The value of RETURNED\_SQLSTATE is the SQLSTATE parameter that would have been returned if this were the only completion or exception condition possible.
  - e) If the value of RETURNED\_SQLSTATE corresponds to *warning* with a subclass of *cursor operation conflict*, then the value of CURSOR\_NAME is the name of the cursor that caused the completion condition to be raised.
  - f) If the value of RETURNED\_SQLSTATE corresponds to *integrity constraint violation*, *transaction rollback — integrity constraint violation*, or a *triggered data change violation* that was caused by a violation of a referential constraint, then:
    - i) The values of CONSTRAINT\_CATALOG and CONSTRAINT\_SCHEMA are the <catalog name> and the <unqualified schema name> of the <schema name> of the schema containing the constraint or assertion. The value of CONSTRAINT\_NAME is the <qualified identifier> of the constraint or assertion.
    - ii) Case:
      - 1) If the violated integrity constraint is a table constraint, then the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are the <catalog name>, the <unqualified schema name> of the <schema name>, and the <qualified identifier>, respectively, of the table in which the table constraint is contained.
      - 2) If the violated integrity constraint is an assertion and if only one table referenced by the assertion has been modified as a result of executing the SQL-statement, then the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are the <catalog name>, the <unqualified schema name> of the <schema name>, and the <qualified identifier>, respectively, of the modified table.

- 3) Otherwise, the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are a zero-length string.

If TABLE\_NAME identifies a declared local temporary table, then CATALOG\_NAME is a zero-length string and SCHEMA\_NAME is “MODULE”.

- g) If the value of RETURNED\_SQLSTATE corresponds to *triggered action exception*, *transaction rollback — triggered action exception*, or a *triggered data change violation* that was caused by a trigger, then:
- i) The values of TRIGGER\_CATALOG and TRIGGER\_SCHEMA are the <catalog name> and the <unqualified schema name> of the <schema name> of the schema containing the trigger. The value of TRIGGER\_NAME is the <qualified identifier> of the <trigger name> of the trigger.
  - ii) The values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are the <catalog name>, the <unqualified schema name> of the <schema name>, and the <qualified identifier> of the <table name>, respectively, of the table on which the trigger is defined.
- h) If the value of RETURNED\_SQLSTATE corresponds to *syntax error or access rule violation*, then:
- i) Case:
    - 1) If the syntax error or access rule violation was caused by reference to a specific table, then the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are  
Case:
      - A) If the specific table referenced was not a declared local temporary table, then the <catalog name>, the <unqualified schema name> of the <schema name> of the schema that contains the table that caused the syntax error or access rule violation, and the <qualified identifier>, respectively.
      - B) Otherwise, the zero-length string, “MODULE”, and the <qualified identifier>, respectively.
    - 2) Otherwise, CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME contain a zero-length string.
  - ii) If the syntax error or access rule violation was for an inaccessible column, then the value of COLUMN\_NAME is the <column name> of that column. Otherwise, the value of COLUMN\_NAME is a zero-length string.
- i) If the value of RETURNED\_SQLSTATE corresponds to *invalid cursor state*, then the value of CURSOR\_NAME is the name of the cursor that is in the invalid state.
- j) If the value of RETURNED\_SQLSTATE corresponds to *with check option violation*, then the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are the <catalog name>, the <unqualified schema name> of the <schema name> of the schema that contains the view that caused the violation of the WITH CHECK OPTION, and the <qualified identifier> of that view, respectively.
- k) If the value of RETURNED\_SQLSTATE does not correspond to *syntax error or access rule violation*, then:
- i) If the values of CATALOG\_NAME, SCHEMA\_NAME, TABLE\_NAME, and COLUMN\_NAME identify a column for which no privileges are granted to the enabled authorization identifiers, then the value of COLUMN\_NAME is replaced by a zero-length string.

## 23.1 &lt;get diagnostics statement&gt;

- ii) If the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME identify a table for which no privileges are granted to the enabled authorization identifiers, then the values of CATALOG\_NAME, SCHEMA\_NAME, and TABLE\_NAME are replaced by a zero-length string.
  - iii) If the values of CONSTRAINT\_CATALOG, CONSTRAINT\_SCHEMA, and CONSTRAINT\_NAME identify a <table constraint> for some table *T* and if no privileges for *T* are granted to the enabled authorization identifiers, then the values of CONSTRAINT\_CATALOG, CONSTRAINT\_SCHEMA, and CONSTRAINT\_NAME are replaced by a zero-length string.
  - iv) If the values of CONSTRAINT\_CATALOG, CONSTRAINT\_SCHEMA, and CONSTRAINT\_NAME identify an assertion contained in some schema *S* and if the owner of *S* is not included in the set of enabled authorization identifiers, then the values of CONSTRAINT\_CATALOG, CONSTRAINT\_SCHEMA, and CONSTRAINT\_NAME are replaced by a zero-length string.
- l) If the value of RETURNED\_SQLSTATE corresponds to *external routine invocation exception*, *external routine exception*, *SQL routine exception*, or *warning*, then
- i) The values of ROUTINE\_CATALOG and ROUTINE\_SCHEMA are the <catalog name> and the <unqualified schema name>, respectively, of the <schema name> of the schema containing the SQL-invoked routine.
  - ii) The values of ROUTINE\_NAME and SPECIFIC\_NAME are the <identifier> of the <routine name> and the <identifier> of the <specific name> of the SQL-invoked routine, respectively.
  - iii) Case:
    - 1) If the condition is related to parameter  $P_i$  of the SQL-invoked routine, then:
      - A) The value of PARAMETER\_MODE is the <parameter mode> of  $P_i$ .
      - B) The value of PARAMETER\_ORDINAL\_POSITION is the value of  $i$ .
      - C) The value of PARAMETER\_NAME is a zero-length string.
    - 2) Otherwise:
      - A) The value of PARAMETER\_MODE is a zero-length string.
      - B) The value of PARAMETER\_ORDINAL\_POSITION is 0 (zero).
      - C) The value of PARAMETER\_NAME is a zero-length string.
- m) If the value of RETURNED\_SQLSTATE corresponds to *external routine invocation exception*, *external routine exception*, *SQL routine exception*, or *warning*, then the value of MESSAGE\_TEXT is the message text item of the SQL-invoked routine that raised the exception. Otherwise the value of MESSAGE\_TEXT is an implementation-defined character string.
- NOTE 479 — An SQL-implementation may set this to <space>s, to a zero-length string, or to a character string describing the condition indicated by RETURNED\_SQLSTATE.
- n) 04 The value of MESSAGE\_LENGTH is the length in characters of the character string value in MESSAGE\_TEXT.
  - o) The value of MESSAGE\_OCTET\_LENGTH is the length in octets of the character string value in MESSAGE\_TEXT.

- p) 04 The values of CONNECTION\_NAME and SERVER\_NAME are respectively

Case:

- i) If COMMAND\_FUNCTION or DYNAMIC\_FUNCTION identifies an <SQL connection statement>, then the <connection name> and the <SQL-server name> specified by or implied by the <SQL connection statement>.
  - ii) Otherwise, the <connection name> and <SQL-server name> of the SQL-session in which the condition was raised.
- q) If the value of RETURNED\_SQLSTATE corresponds to *data exception — numeric value out of range*, *data exception — invalid character value for cast*, *data exception — string data, right truncation*, *data exception — interval field overflow*, *integrity constraint violation*, or *warning — string data, right truncation*, and the condition was raised as the result of an assignment to an SQL parameter during an SQL-invoked routine invocation, then:
- i) The values of ROUTINE\_CATALOG and ROUTINE\_SCHEMA are the <catalog name> and the <unqualified schema name>, respectively, of the <schema name> of the schema containing the routine.
  - ii) The values of the ROUTINE\_NAME and SPECIFIC\_NAME are the <identifier> of the <routine name> and the <identifier> of the <specific name>, respectively, of the routine.
  - iii) If the condition is related to parameter  $P_i$  of the SQL-invoked routine, then:
    - 1) The value of PARAMETER\_MODE is the <parameter mode> of  $P_i$ .
    - 2) The value of PARAMETER\_ORDINAL\_POSITION is the value of  $i$ .
    - 3) If an <SQL parameter name> was specified for the SQL parameter when the SQL-invoked routine was created, then the value of PARAMETER\_NAME is the <SQL parameter name> of  $P_i$ . Otherwise, the value of PARAMETER\_NAME is a zero-length string.
- 5) The values of character string items where not otherwise specified by the preceding rules are set to a zero-length string.

NOTE 480 — There are no numeric items that are not set by these rules.

- 6) Case:

- a) If <SQL diagnostics information> contains an <all information>, then let  $N$  be the number of occupied condition areas in  $DA$ . Let  $S$  be the <literal> ' '. For each  $X$  such that  $X$  is either the <statement information item name> of a statement information item in  $DA$ , or the <condition information item name> of a condition information item in an occupied condition area of  $DA$ , or CONDITION\_NUMBER (whose value is the ordinal position of an occupied condition area), let  $CHARFORM ( X )$  denote the <value expression>

$CAST ( LX AS CHARACTER VARYING ( M ) )$

where  $LX$  is a <literal> denoting the value of  $X$  and  $M$  is the implementation-defined maximum length of variable-length character strings. Let  $PRESENTATION ( X )$  denote the <value expression>

'X' = ' || CHARFORM ( X )

Let SEMICOLONCAT (  $str, din$  ) be equivalent to

## 23.1 &lt;get diagnostics statement&gt;

```

CASE
 WHEN CHARFORM (din) = '' THEN str ;
 WHEN str = '' THEN PRESENTATION (din) ;
 ELSE str || ';' || PRESENTATION (din) ;
END CASE

```

- i) If <all qualifier> is omitted or contains STATEMENT, then for each statement information item *SII* in *DA*, taken in an implementation-dependent order, let *S* be

*SEMICOLONCAT* ( *S* , *SIN* )

where *SIN* is the <statement information item name> of *SII*.

- ii) If <all qualifier> is omitted or contains CONDITION and does not contain a <condition number>, then for each *CN*,  $1 \text{ (one)} \leq CN \leq N$ , in an implementation-defined order:

- 1) Let *S* be

*SEMICOLONCAT* ( *S* , *COND* )

where *COND* is *CONDITION\_NUMBER* and the value of *CONDITION\_NUMBER* is *CN*.

- 2) For each condition information item *CII* in condition area *CN*, taken in an implementation-dependent order, let *S* be

*SEMICOLONCAT* ( *S* , *CIN* )

where *CIN* is the <condition number> of *CII*.

- iii) If <all qualifier> contains CONDITION and a <condition number>, then let *CN* be the value of <condition number>. For each condition information item *CII* in condition area *CN*, taken in an implementation-dependent order, let *S* be

*SEMICOLONCAT* ( *S* , *CIN* )

where *CIN* is the <condition number> of *CII*.

- iv) The General Rules of [Subclause 9.1, “Retrieval assignment”](#), apply to <all info target> and *S*, as *TARGET* and *VALUE*, respectively.

- b) If <statement information> is specified, then for each <statement information item>, the General Rules of [Subclause 9.1, “Retrieval assignment”](#), are applied with <simple target specification> as *TARGET* and the value of <statement information item name> as *VALUE*.

- c) Otherwise (<condition information> is specified), for each <condition information item>, the General Rules of [Subclause 9.1, “Retrieval assignment”](#), are applied with <simple target specification> as *TARGET* and the value of <condition information item name> as *VALUE*.

## Conformance Rules

- 1) Without Feature F121, “Basic diagnostics management”, conforming SQL language shall not contain a <get diagnostics statement>.

**23.1 <get diagnostics statement>**

- 2) Without Feature T511, “Transaction counts”, conforming SQL language shall not contain a <statement information item name> that contains TRANSACTIONS\_COMMITTED, TRANSACTIONS\_ROLLED\_BACK, or TRANSACTION\_ACTIVE.
- 3) Without Feature F122, “Enhanced diagnostics management”, conforming SQL language shall not contain a <get diagnostics statement> containing a <simple target specification> whose declared type is different from that of the corresponding <statement information item name> or <condition information item name>.
- 4) Without Feature F123, “All diagnostics”, conforming SQL language shall not contain an <all information>.

## 23.2 Pushing and popping the diagnostics area stack

### Function

Define operations on the diagnostics area stack.

### Syntax Rules

*None.*

### Access Rules

*None.*

### General Rules

- 1) Let *OP* be the *OPERATION* and let *DAS* be the *STACK* specified in an application of this Subclause.
- 2) Case:
  - a) If *OP* is “PUSH”, then

Case:

    - i) If the number of diagnostics areas in *DAS* is equal to the implementation-dependent maximum number of diagnostics areas per diagnostics area stack, then an exception condition is raised: *diagnostics exception — maximum number of stacked diagnostics areas exceeded*.
    - ii) Otherwise, *DAS* is pushed and the contents of the second diagnostics area in *DAS* are copied to the first.
  - b) If *OP* is “POP”, then the first diagnostics area is removed from *DAS* such that all subsequent diagnostics areas in *DAS* move up one position, the second becoming the first, the third becoming the second, and so on.

### Conformance Rules

*None.*



## 24 Status codes

### 24.1 SQLSTATE

*This Subclause is modified by Subclause 20.1, “SQLSTATE”, in ISO/IEC 9075-4.*

*This Subclause is modified by Subclause 26.1, “SQLSTATE”, in ISO/IEC 9075-9.*

*This Subclause is modified by Subclause 15.1, “SQLSTATE”, in ISO/IEC 9075-10.*

*This Subclause is modified by Subclause 15.1, “SQLSTATE”, in ISO/IEC 9075-13.*

*This Subclause is modified by Subclause 23.1, “SQLSTATE”, in ISO/IEC 9075-14.*

The character string value returned in an SQLSTATE parameter comprises a 2-character class value followed by a 3-character subclass value, each with an implementation-defined character set that has a one-octet character encoding form and is restricted to <digit>s and <simple Latin upper case letter>s. Table 33, “SQLSTATE class and subclass values”, specifies the class value for each condition and the subclass value or values for each class value.

Class values that begin with one of the <digit>s '0', '1', '2', '3', or '4' or one of the <simple Latin upper case letter>s 'A', 'B', 'C', 'D', 'E', 'F', 'G', or 'H' are returned only for conditions defined in ISO/IEC 9075 or in any other International Standard. The range of such class values are called *standard-defined classes*. Some such class codes are reserved for use by specific International Standards, as specified elsewhere in this Clause. Subclass values associated with such classes that also begin with one of those 13 characters are returned only for conditions defined in ISO/IEC 9075 or some other International Standard. The range of such class values are called *standard-defined classes*. Subclass values associated with such classes that begin with one of the <digit>s '5', '6', '7', '8', or '9' or one of the <simple Latin upper case letter>s 'I', 'J', 'K', 'L', 'M', 'N', 'O', 'P', 'Q', 'R', 'S', 'T', 'U', 'V', 'W', 'X', 'Y', or 'Z' are reserved for implementation-specified conditions and are called *implementation-defined subclasses*.

Class values that begin with one of the <digit>s '5', '6', '7', '8', or '9' or one of the <simple Latin upper case letter>s 'I', 'J', 'K', 'L', 'M', 'N', 'O', 'P', 'Q', 'R', 'S', 'T', 'U', 'V', 'W', 'X', 'Y', or 'Z' are reserved for implementation-specified exception conditions and are called *implementation-defined classes*. All subclass values except '000', which means *no subclass*, associated with such classes are reserved for implementation-specified conditions and are called *implementation-defined subclasses*. An implementation-defined completion condition shall be indicated by returning an implementation-defined subclass in conjunction with one of the classes *successful completion*, *warning*, or *no data*.

If a subclass value is not specified for a condition, then either subclass '000' or an implementation-defined subclass is returned.

NOTE 481 — One consequence of this is that an SQL-implementation may, but is not required by ISO/IEC 9075 to, provide subcodes for exception condition *syntax error or access rule violation* that distinguish between the syntax error and access rule violation cases.

If multiple completion conditions: *warning* or multiple exception conditions, including implementation-defined exception conditions, are raised, then it is implementation-dependent which of the corresponding SQLSTATE values is returned in the SQLSTATE status parameter, provided that the precedence rules in Subclause 4.29.2, “Status parameters”, are obeyed. Any number of applicable conditions values in addition to the one returned in the SQLSTATE status parameter, may be returned in the diagnostics area.

## 24.1 SQLSTATE

An implementation-specified condition may duplicate, in whole or in part, a condition defined in ISO/IEC 9075; however, if such a condition occurs as a result of executing a statement, then the corresponding implementation-defined SQLSTATE value shall not be returned in the SQLSTATE parameter but may be returned in the diagnostics area.

The “Category” column has the following meanings: “S” means that the class value given corresponds to successful completion and is a completion condition; “W” means that the class value given corresponds to a successful completion but with a warning and is a completion condition; “N” means that the class value given corresponds to a no-data situation and is a completion condition; “X” means that the class value given corresponds to an exception condition.

**Table 33 — SQLSTATE class and subclass values**

Category	Condition	Class	Subcondition	Subclass
X	<i>ambiguous cursor name</i>	3C	<i>(no subclass)</i>	000
X	<i>attempt to assign to non-updatable column</i>	0U	<i>(no subclass)</i>	000
X	<i>attempt to assign to ordering column</i>	0V	<i>(no subclass)</i>	000
X	<i>cardinality violation</i>	21	<i>(no subclass)</i>	000
X	<i>connection exception</i>	08	<i>(no subclass)</i>	000
			<i>connection does not exist</i>	003
			<i>connection failure</i>	006
			<i>connection name in use</i>	002
			<i>SQL-client unable to establish SQL-connection</i>	001
			<i>SQL-server rejected establishment of SQL-connection</i>	004
			<i>transaction resolution unknown</i>	007
X	<i>cursor sensitivity exception</i>	36	<i>(no subclass)</i>	000
			<i>request failed</i>	002
			<i>request rejected</i>	001
X	<i>data exception</i>	22	<i>(no subclass)</i>	000
			<i>array data, right truncation</i>	02F

Category	Condition	Class	Subcondition	Subclass
			<i>array element error</i>	02E
			<i>attempt to replace a zero-length string</i>	01U
			<i>character not in repertoire</i>	021
			<i>datetime field overflow</i>	008
			<i>division by zero</i>	012
			<i>error in assignment</i>	005
			<i>escape character conflict</i>	00B
			<i>indicator overflow</i>	022
			<i>interval field overflow</i>	015
			<i>interval value out of range</i>	00P
			<i>invalid argument for natural logarithm</i>	01E
			<i>invalid argument for power function</i>	01F
			<i>invalid argument for width bucket function</i>	01G
			<i>invalid character value for cast</i>	018
			<i>invalid datetime format</i>	007
			<i>invalid escape character</i>	019
			<i>invalid escape octet</i>	00D
			<i>invalid escape sequence</i>	025
			<i>invalid indicator parameter value</i>	010
			<i>invalid interval format</i>	006
			<i>invalid parameter value</i>	023
			<i>invalid preceding or following size in window function</i>	013
			<i>invalid regular expression</i>	01B

Category	Condition	Class	Subcondition	Subclass
			<i>invalid repeat argument in a sample clause</i>	02G
			<i>invalid sample size</i>	02H
			<i>invalid time zone displacement value</i>	009
			<i>invalid use of escape character</i>	00C
			<i>invalid XQuery option flag</i>	01T
			<i>invalid XQuery regular expression</i>	01S
			<i>invalid XQuery replacement string</i>	01V
			<i>most specific type mismatch</i>	00G
			<i>multiset value overflow</i>	00Q
			<i>noncharacter in UCS string</i>	029
			<i>null value substituted for mutator subject parameter</i>	02D
			<i>null row not permitted in table</i>	01C
			<i>null value in array target</i>	00E
			<i>null value, no indicator parameter</i>	002
			<i>null value not allowed</i>	004
			<i>numeric value out of range</i>	003
			<i>sequence generator limit exceeded</i>	00H
			<i>string data, length mismatch</i>	026
			<i>string data, right truncation</i>	001
			<i>substring error</i>	011
			<i>trim error</i>	027
			<i>unterminated C string</i>	024
			<i>zero-length character string</i>	00F

Category	Condition	Class	Subcondition	Subclass
X	<i>dependent privilege descriptors still exist</i>	2B	(no subclass)	000
X	<i>diagnostics exception</i>	0Z	(no subclass)	000
			<i>maximum number of stacked diagnostics areas exceeded</i>	001
X	<i>dynamic SQL error</i>	07	(no subclass)	000
			<i>cursor specification cannot be executed</i>	003
			<i>data type transform function violation</i>	00B
			<i>invalid DATA target</i>	00D
			<i>invalid DATETIME_INTERVAL_CODE</i>	00F
			<i>invalid descriptor count</i>	008
			<i>invalid descriptor index</i>	009
			<i>invalid LEVEL value</i>	00E
			<i>prepared statement not a cursor specification</i>	005
			<i>restricted data type attribute violation</i>	006
			<i>undefined DATA value</i>	00C
			<i>using clause does not match dynamic parameter specifications</i>	001
			<i>using clause does not match target specifications</i>	002
			<i>using clause required for dynamic parameters</i>	004
			<i>using clause required for result fields</i>	007
X	<i>external routine exception</i>	38	(no subclass)	000
			<i>containing SQL not permitted</i>	001

## 24.1 SQLSTATE

Category	Condition	Class	Subcondition	Subclass
			<i>modifying SQL-data not permitted</i>	002
			<i>prohibited SQL-statement attempted</i>	003
			<i>reading SQL-data not permitted</i>	004
X	<i>external routine invocation exception</i>	39	(no subclass)	000
			<i>null value not allowed</i>	004
X	<i>feature not supported</i>	0A	(no subclass)	000
			<i>multiple server transactions</i>	001
X	<i>integrity constraint violation</i>	23	(no subclass)	000
			<i>restrict violation</i>	001
X	<i>invalid authorization specification</i>	28	(no subclass)	000
X	<i>invalid catalog name</i>	3D	(no subclass)	000
X	<i>invalid character set name</i>	2C	(no subclass)	000
X	<i>invalid condition number</i>	35	(no subclass)	000
X	<i>invalid connection name</i>	2E	(no subclass)	000
X	<i>invalid cursor name</i>	34	(no subclass)	000
X	<i>invalid cursor state</i>	24	(no subclass)	000
X	<i>invalid grantor</i>	0L	(no subclass)	000
X	<i>invalid role specification</i>	0P	(no subclass)	000
X	<i>invalid schema name</i>	3F	(no subclass)	000
X	<i>invalid schema name list specification</i>	0E	(no subclass)	000
X	<i>invalid collation name</i>	2H	(no subclass)	000
X	<i>invalid SQL descriptor name</i>	33	(no subclass)	000
X	<i>invalid SQL-invoked procedure reference</i>	0M	(no subclass)	000

Category	Condition	Class	Subcondition	Subclass
X	<i>invalid SQL statement name</i>	26	(no subclass)	000
X	<i>invalid SQL statement identifier</i>	30	(no subclass)	000
X	<i>invalid target type specification</i>	0D	(no subclass)	000
X	<i>invalid transaction state</i>	25	(no subclass)	000
			<i>active SQL-transaction</i>	001
			<i>branch transaction already active</i>	002
			<i>held cursor requires same isolation level</i>	008
			<i>inappropriate access mode for branch transaction</i>	003
			<i>inappropriate isolation level for branch transaction</i>	004
			<i>no active SQL-transaction for branch transaction</i>	005
			<i>read-only SQL-transaction</i>	006
			<i>schema and data statement mixing not supported</i>	007
X	<i>invalid transaction termination</i>	2D	(no subclass)	000
X	<i>invalid transform group name specification</i>	0S	(no subclass)	000
X	<i>locator exception</i>	0F	(no subclass)	000
			<i>invalid specification</i>	001
N	<i>no data</i>	02	(no subclass)	000
			<i>no additional result sets returned</i>	001
X	<i>prohibited statement encountered during trigger execution</i>	0W	(no subclass)	000
X	Remote Database Access	HZ	(See Table 34, “SQLSTATE class codes for RDA”, for the definition of protocol subconditions and subclass code values)	

Category	Condition	Class	Subcondition	Subclass
X	<i>savepoint exception</i>	3B	(no subclass)	000
			<i>invalid specification</i>	001
			<i>too many</i>	002
X	<i>SQL routine exception</i>	2F	(no subclass)	000
			<i>function executed no return state- ment</i>	005
			<i>modifying SQL-data not permitted</i>	002
			<i>prohibited SQL-statement attempted</i>	003
			<i>reading SQL-data not permitted</i>	004
S	<i>successful completion</i>	00	(no subclass)	000
X	<i>syntax error or access rule viola- tion</i>	42	(no subclass)	000
X	<i>target table disagrees with cursor specification</i>	0T	(no subclass)	000
X	<i>transaction rollback</i>	40	(no subclass)	000
			<i>integrity constraint violation</i>	002
			<i>serialization failure</i>	001
			<i>statement completion unknown</i>	003
			<i>triggered action exception</i>	004
X	<i>triggered action exception</i>	09	(no subclass)	000
X	<i>triggered data change violation</i>	27	(no subclass)	000
W	<i>warning</i>	01	(no subclass)	000
			<i>additional result sets returned</i>	00D
			<i>array data, right truncation</i>	02F
			<i>attempt to return too many result sets</i>	00E
			<i>cursor operation conflict</i>	001



Category	Condition	Class	Subcondition	Subclass
			<i>default value too long for information schema</i>	00B
			<i>disconnect error</i>	002
			<i>insufficient item descriptor areas</i>	005
			<i>invalid number of conditions</i>	012
			<i>null value eliminated in set function</i>	003
			<i>privilege not granted</i>	007
			<i>privilege not revoked</i>	006
			<i>query expression too long for information schema</i>	00A
			<i>result sets returned</i>	00C
			<i>search condition too long for information schema</i>	009
			<i>statement too long for information schema</i>	00F
			<i>string data, right truncation</i>	004
X	<i>with check option violation</i>	44	<i>(no subclass)</i>	000

## 24.2 Remote Database Access SQLSTATE Subclasses

ISO/IEC 9075 reserves SQLSTATE class 'HZ' for Remote Database Access errors, which may occur when an SQL-client interacts with an SQL-server across a communications network using an RDA Application Context. [ISO9579], [ISO8649], and [ISO10026] define a number of exception conditions that shall be detected in a conforming ISO RDA implementation. This Subclause defines SQLSTATE subclass codes for each such condition out of the set of codes reserved for International Standards.

If an implementation using RDA reports a condition shown in Table 34, “SQLSTATE class codes for RDA”, for a given exception condition, then it shall use the SQLSTATE class code 'HZ' and the subclass codes shown, and shall set the values of CLASS\_ORIGIN to 'ISO 9075' and SUBCLASS\_ORIGIN as indicated in Table 34, “SQLSTATE class codes for RDA”, when those exceptions are retrieved by a <get diagnostics statement>.

An implementation using client-server communications other than RDA may report conditions corresponding to the conditions shown in Table 34, “SQLSTATE class codes for RDA”, using the SQLSTATE class code 'HZ' and the corresponding subclass codes shown. It may set the values of CLASS\_ORIGIN to 'ISO 9075' and SUBCLASS\_ORIGIN as indicated in Table 34, “SQLSTATE class codes for RDA”. Any other communications error shall be returned with a subclass code from the implementation-defined range, with CLASS\_ORIGIN set to 'ISO 9075' and SUBCLASS\_ORIGIN set to an implementation-defined character string.

A Remote Database Access exception may also result in an SQL completion condition defined in Table 33, “SQLSTATE class and subclass values” (such as '40000', *transaction rollback*); if such a condition occurs, then the 'HZ' class SQLSTATE shall not be returned in the SQLSTATE parameter, but may be returned in the Diagnostics Area.

**Table 34 — SQLSTATE class codes for RDA**

SQLSTATE Class	Subclass Origin
HZ	[ISO9579]

## 25 Conformance

### 25.1 Claims of conformance to SQL/Foundation

In addition to the requirements of [ISO9075-1], in [Clause 8, “Conformance”](#), a claim of conformance to this part of ISO/IEC 9075 shall:

- 1) Claim conformance to at least one of:
  - Feature B011, “Embedded Ada”
  - Feature B012, “Embedded C”
  - Feature B013, “Embedded COBOL”
  - Feature B014, “Embedded Fortran”
  - Feature B015, “Embedded MUMPS”
  - Feature B016, “Embedded Pascal”
  - Feature B017, “Embedded PL/I”
  - Feature B111, “Module language Ada”
  - Feature B112, “Module language C”
  - Feature B113, “Module language COBOL”
  - Feature B114, “Module language Fortran”
  - Feature B115, “Module language MUMPS”
  - Feature B116, “Module language Pascal”
  - Feature B117, “Module language PL/I”
- 2) A claim of conformance to at least one of:
  - Feature B121, “Routine language Ada”
  - Feature B122, “Routine language C”
  - Feature B123, “Routine language COBOL”
  - Feature B124, “Routine language Fortran”
  - Feature B125, “Routine language MUMPS”
  - Feature B126, “Routine language Pascal”
  - Feature B127, “Routine language PL/I”

**25.1 Claims of conformance to SQL/Foundation**

— Feature B128, “Routine language SQL”

**25.2 Additional conformance requirements for SQL/Foundation**

An SQL-implementation that claims conformance to a feature in this part of ISO/IEC 9075 shall also claim conformance to the same feature, if present, in [ISO9075-11].

An SQL-implementation that claims conformance to Feature T061, “UCS support”, shall:

- Conform to [ISO10646] at some specified level.
- Provide at least one of the named character sets UTF8, UTF16, and UTF32.
- Provide, as the default collation for each such character set, a collation that conforms to [ISO14651] at some level.

**25.3 Implied feature relationships of SQL/Foundation**

**Table 35 — Implied feature relationships of SQL/Foundation**

<b>Feature ID</b>	<b>Feature Name</b>	<b>Implied Feature ID</b>	<b>Implied Feature Name</b>
B032	Extended dynamic SQL	B031	Basic dynamic SQL
B034	Dynamic specification of cursor attributes	B031	Basic dynamic SQL
B035	Non-extended descriptor names	B031	Basic dynamic SQL
B111	Module language Ada	E182	Module language
B112	Module language C	E182	Module language
B113	Module language COBOL	E182	Module language
B114	Module language Fortran	E182	Module language
B115	Module language MUMPS	E182	Module language
B116	Module language Pascal	E182	Module language
B117	Module language PL/I	E182	Module language
F053	OVERLAPS predicate	F052	Intervals and date arithmetic
F313	Enhanced MERGE statement	F312	MERGE statement

## 25.3 Implied feature relationships of SQL/Foundation

Feature ID	Feature Name	Implied Feature ID	Implied Feature Name
F381	Extended schema manipulation	F491	Constraint management
F451	Character set definition	F461	Named character sets
F521	Assertions	F491	Constraint management
F693	SQL-session and client module collations	F690	Collation support
F711	ALTER domain	F251	Domain support
F721	Deferrable constraints	F491	Constraint management
F762	CURRENT_CATALOG	F651	Catalog name qualifiers
F762	CURRENT_CATALOG	F763	CURRENT_SCHEMA
F801	Full set function	F441	Extended set function support
F846	Octet support in regular expression operators	T061	UCS support
F851	<order by clause> in subqueries	F850	Top-level <order by clause> in <query expression>
F855	Nested <order by clause> in <query expression>	F850	Top-level <order by clause> in <query expression>
F856	Nested <fetch first clause> in <query expression>	F857	Top-level <fetch first clause> in <query expression>
S024	Enhanced structured types	S023	Basic structured types
S041	Basic reference types	S051	Create table of type
S043	Enhanced reference types	S041	Basic reference types
S051	Create table of type	S023	Basic structured types
S081	Subtables	S051	Create table of type
S092	Arrays of user-defined types	S091	Basic array support
S094	Arrays of reference types	S041	Basic reference types
S094	Arrays of reference types	S091	Basic array support
S095	Array constructors by query	S091	Basic array support

## 25.3 Implied feature relationships of SQL/Foundation

Feature ID	Feature Name	Implied Feature ID	Implied Feature Name
S096	Optional array bounds	S091	Basic array support
S111	ONLY in query expressions	S051	Create table of type
S201	SQL-invoked routines on arrays	S091	Basic array support
S202	SQL-invoked routines on multisets	S271	Basic multiset support
S231	Structured type locators	S023	Basic structured types
S232	Array locators	S091	Basic array support
S233	Multiset locators	S271	Basic multiset support
S242	Alter transform statement	S241	Transform functions
S272	Multisets of user-defined types	S271	Basic multiset support
S274	Multisets of reference types	S041	Basic reference types
S274	Multisets of reference types	S271	Basic multiset support
S275	Advanced multiset support	S271	Basic multiset support
T022	Advanced support for BINARY and VARBINARY data types	T021	BINARY and VARBINARY data types
T042	Extended LOB data type support	T041	Basic LOB data type support
T043	Multiplier T	T041	Basic LOB data type support
T044	Multiplier P	T043	Multiplier T
T061	UCS Support	F461	Named character sets
T061	UCS support	F690	Collation support
T122	WITH (excluding RECURSIVE) in subquery	T121	WITH (excluding RECURSIVE) in query expression
T131	Recursive query	T121	WITH (excluding RECURSIVE) in query expression
T132	Recursive query in subquery	T122	WITH (excluding RECURSIVE) in subquery
T132	Recursive query in subquery	T131	Recursive query

## 25.3 Implied feature relationships of SQL/Foundation

Feature ID	Feature Name	Implied Feature ID	Implied Feature Name
T152	DISTINCT predicate with negation	T151	DISTINT predicate
T173	Extended LIKE clause in table definition	T171	LIKE clause in table definition
T212	Enhanced trigger capability	T211	Basic trigger capability
T213	INSTEAD OF triggers	T211	Basic trigger capability
T272	Enhanced savepoint management	T271	Savepoints
T332	Extended roles	T331	Basic roles
T432	Nested and concatenated GROUPING SETS	T431	Extended grouping capabilities
T433	Multiargument GROUPING function	T431	Extended grouping capabilities
T511	Transaction counts	F121	Basic diagnostics management
T571	Array-returning external SQL-invoked functions	S201	SQL-invoked routines on arrays
T572	Multiset-returning external SQL-invoked functions	S202	SQL-invoked routines on multisets
T612	Advanced OLAP operations	T611	Elementary OLAP operations
T652	SQL-dynamic statements in SQL routines	B031	Basic dynamic SQL
T654	SQL-dynamic statements in external routines	B031	Basic dynamic SQL

*(Blank page)*



## Annex A (informative)

### SQL Conformance Summary

*This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-3.  
 This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-4.  
 This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-9.  
 This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-10.  
 This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-11.  
 This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-13.  
 This Annex is modified by Annex A, “SQL Conformance Summary”, in ISO/IEC 9075-14.*

The contents of this Annex summarizes all Conformance Rules, ordered by Feature ID and by Subclause.

- 1) Specifications for Feature B011, “Embedded Ada”:
  - a) Subclause 21.3, “<embedded SQL Ada program>”:
    - i) Without Feature B011, “Embedded Ada”, conforming SQL language shall not contain an <embedded SQL Ada program>.
- 2) Specifications for Feature B012, “Embedded C”:
  - a) Subclause 21.4, “<embedded SQL C program>”:
    - i) Without Feature B012, “Embedded C”, conforming SQL language shall not contain an <embedded SQL C program>.
- 3) Specifications for Feature B013, “Embedded COBOL”:
  - a) Subclause 21.5, “<embedded SQL COBOL program>”:
    - i) Without Feature B013, “Embedded COBOL”, conforming SQL language shall not contain an <embedded SQL COBOL program>.
- 4) Specifications for Feature B014, “Embedded Fortran”:
  - a) Subclause 21.6, “<embedded SQL Fortran program>”:
    - i) Without Feature B014, “Embedded Fortran”, conforming SQL language shall not contain an <embedded SQL Fortran program>.
- 5) Specifications for Feature B015, “Embedded MUMPS”:
  - a) Subclause 21.7, “<embedded SQL MUMPS program>”:
    - i) Without Feature B015, “Embedded MUMPS”, conforming SQL language shall not contain an <embedded SQL MUMPS program>.
- 6) Specifications for Feature B016, “Embedded Pascal”:

- a) Subclause 21.8, “<embedded SQL Pascal program>”:
  - i) Without Feature B016, “Embedded Pascal”, conforming SQL language shall not contain an <embedded SQL Pascal program>.
- 7) Specifications for Feature B017, “Embedded PL/I”:
  - a) Subclause 21.9, “<embedded SQL PL/I program>”:
    - i) Without Feature B017, “Embedded PL/I”, conforming SQL language shall not contain an <embedded SQL PL/I program>.
- 8) Specifications for Feature B021, “Direct SQL”:
  - a) Subclause 22.1, “<direct SQL statement>”:
    - i) Without Feature B021, “Direct SQL”, conforming SQL language shall not contain a <direct SQL statement>.
- 9) Specifications for Feature B031, “Basic dynamic SQL”:
  - a) Subclause 5.4, “Names and identifiers”:
    - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <SQL statement name>.
    - ii) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain <dynamic cursor name>.
    - iii) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <descriptor name>.
  - b) Subclause 6.4, “<value specification> and <target specification>”:
    - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <general value specification> that contains a <dynamic parameter specification>.
  - c) Subclause 20.2, “<allocate descriptor statement>”:
    - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <allocate descriptor statement>.
  - d) Subclause 20.3, “<deallocate descriptor statement>”:
    - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <deallocate descriptor statement>.
  - e) Subclause 20.4, “<get descriptor statement>”:
    - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <get descriptor statement>.
  - f) Subclause 20.5, “<set descriptor statement>”:
    - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <set descriptor statement>.
  - g) Subclause 20.6, “<prepare statement>”:

- i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <prepare statement>.
- h) Subclause 20.9, “<describe statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <describe output statement>.
- i) Subclause 20.10, “<input using clause>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <input using clause>.
- j) Subclause 20.11, “<output using clause>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <output using clause>.
- k) Subclause 20.12, “<execute statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <execute statement>.
- l) Subclause 20.13, “<execute immediate statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain an <execute immediate statement>.
- m) Subclause 20.14, “<dynamic declare cursor>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic declare cursor>.
- n) Subclause 20.16, “<dynamic open statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic open statement>.
- o) Subclause 20.17, “<dynamic fetch statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic fetch statement>.
- p) Subclause 20.18, “<dynamic single row select statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic single row select statement>.
- q) Subclause 20.19, “<dynamic close statement>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic close statement>.
- r) Subclause 20.20, “<dynamic delete statement: positioned>”:
  - i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic delete statement: positioned>.
- s) Subclause 20.21, “<dynamic update statement: positioned>”:

- i) Without Feature B031, “Basic dynamic SQL”, conforming SQL language shall not contain a <dynamic update statement: positioned>.

10) Specifications for Feature B032, “Extended dynamic SQL”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <extended statement name> or <extended cursor name>.
- ii) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <descriptor name> that is not a <literal> or a <non-extended descriptor name>.

b) Subclause 20.2, “<allocate descriptor statement>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain an <occurrences> that is not a <literal>.

c) Subclause 20.8, “<deallocate prepared statement>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <deallocate prepared statement>.

d) Subclause 20.9, “<describe statement>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <describe input statement>.

e) Subclause 20.12, “<execute statement>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <result using clause>.

f) Subclause 20.15, “<allocate cursor statement>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain an <allocate cursor statement>.

g) Subclause 20.22, “<preparable dynamic delete statement: positioned>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <preparable dynamic delete statement: positioned>.

h) Subclause 20.24, “<preparable dynamic update statement: positioned>”:

- i) Without Feature B032, “Extended dynamic SQL”, conforming SQL language shall not contain a <preparable dynamic update statement: positioned>.

11) Specifications for Feature B033, “Untyped SQL-invoked function arguments”:

a) Subclause 10.4, “<routine invocation>”:

- i) Without Feature B033, “Untyped SQL-invoked function arguments”, conforming SQL language shall not contain a <routine invocation> that is not simply contained in a <call statement> that simply contains an <SQL argument> that is a <dynamic parameter specification>.

12) Specifications for Feature B034, “Dynamic specification of cursor attributes”:

a) Subclause 20.6, “<prepare statement>”:

- i) Without Feature B034, “Dynamic specification of cursor attributes”, conforming SQL language shall not contain an <attributes specification>.
- 13) Specifications for Feature B035, “Non-extended descriptor names”:
  - a) Subclause 5.4, “Names and identifiers”:
    - i) Without Feature B035, “Non-extended descriptor names”, conforming SQL language shall not contain a <descriptor name> that is a <non-extended descriptor name>.
- 14) Specifications for Feature B041, “Extensions to embedded SQL exception declarations”:
  - a) Subclause 21.2, “<embedded exception declaration>”:
    - i) Without Feature B041, “Extensions to embedded SQL exception declarations”, conforming SQL language shall not contain an <SQL condition> that contains either SQLSTATE or CONSTRAINT.
- 15) Specifications for Feature B051, “Enhanced execution rights”:
  - a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B051, “Enhanced execution rights”, conforming SQL language shall not contain a <module authorization clause> that immediately contains FOR STATIC ONLY or FOR STATIC AND DYNAMIC.
  - b) Subclause 21.1, “<embedded SQL host program>”:
    - i) Without Feature B051, “Enhanced execution rights”, conforming SQL language shall not contain an <embedded authorization declaration>.
- 16) Specifications for Feature B111, “Module language Ada”:
  - a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B111, “Module language Ada”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains ADA.
- 17) Specifications for Feature B112, “Module language C”:
  - a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B112, “Module language C”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains C.
- 18) Specifications for Feature B113, “Module language COBOL”:
  - a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B113, “Module language COBOL”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains COBOL.
- 19) Specifications for Feature B114, “Module language Fortran”:
  - a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B114, “Module language Fortran”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains FORTRAN.

- 20) Specifications for Feature B115, “Module language MUMPS”:
- a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B115, “Module language MUMPS”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains M.
- 21) Specifications for Feature B116, “Module language Pascal”:
- a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B116, “Module language Pascal”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains PASCAL.
- 22) Specifications for Feature B117, “Module language PL/I”:
- a) Subclause 13.1, “<SQL-client module definition>”:
    - i) Without Feature B117, “Module language PL/I”, conforming SQL language shall not contain an <SQL-client module definition> that contains a <language clause> that contains PLI.
- 23) Specifications for Feature B121, “Routine language Ada”:
- a) Subclause 11.51, “<SQL-invoked routine>”:
    - i) Without Feature B121, “Routine language Ada”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains ADA.
- 24) Specifications for Feature B122, “Routine language C”:
- a) Subclause 11.51, “<SQL-invoked routine>”:
    - i) Without Feature B122, “Routine language C”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains C.
- 25) Specifications for Feature B123, “Routine language COBOL”:
- a) Subclause 11.51, “<SQL-invoked routine>”:
    - i) Without Feature B123, “Routine language COBOL”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains COBOL.
- 26) Specifications for Feature B124, “Routine language Fortran”:
- a) Subclause 11.51, “<SQL-invoked routine>”:
    - i) Without Feature B124, “Routine language Fortran”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains FORTRAN.
- 27) Specifications for Feature B125, “Routine language MUMPS”:
- a) Subclause 11.51, “<SQL-invoked routine>”:
    - i) Without Feature B125, “Routine language MUMPS”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains M.
- 28) Specifications for Feature B126, “Routine language Pascal”:
- a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature B126, “Routine language Pascal”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains PASCAL.

29) Specifications for Feature B127, “Routine language PL/I”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature B127, “Routine language PL/I”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains PLI.

30) Specifications for Feature B128, “Routine language SQL”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature B128, “Routine language SQL”, conforming SQL language shall not contain a <routine characteristic> that contains a <language clause> that contains SQL.

31) Specifications for Feature F032, “CASCADE drop behavior”:

a) Subclause 11.22, “<drop table statement>”:

- i) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop table statement> that contains <drop behavior> that contains CASCADE.

b) Subclause 11.24, “<drop view statement>”:

- i) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop view statement> that contains a <drop behavior> that contains CASCADE.

c) Subclause 11.50, “<drop data type statement>”:

- i) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop data type statement> that contains a <drop behavior> that contains CASCADE.

d) Subclause 11.53, “<drop routine statement>”:

- i) Without Feature F032, “CASCADE drop behavior”, conforming SQL language shall not contain a <drop routine statement> that contains a <drop behavior> that contains CASCADE.

32) Specifications for Feature F033, “ALTER TABLE statement: DROP COLUMN clause”:

a) Subclause 11.19, “<drop column definition>”:

- i) Without Feature F033, “ALTER TABLE statement: DROP COLUMN clause”, conforming SQL language shall not contain a <drop column definition>.

33) Specifications for Feature F034, “Extended REVOKE statement”:

a) Subclause 12.7, “<revoke statement>”:

- i) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke statement> that contains a <drop behavior> that contains CASCADE.
- ii) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke option extension> that contains GRANT OPTION FOR.
- iii) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke statement> that contains a <privileges> that contains an <object name> where



the owner of the SQL-schema that is specified explicitly or implicitly in the <object name> is not the current authorization identifier.

- iv) Without Feature F034, “Extended REVOKE statement”, conforming SQL language shall not contain a <revoke statement> such that there exists a privilege descriptor *PD* that satisfies all the following conditions:
  - 1) *PD* identifies the object identified by <object name> simply contained in <privileges> contained in the <revoke statement>.
  - 2) *PD* identifies the <grantee> identified by any <grantee> simply contained in <revoke statement> and that <grantee> does not identify the owner of the SQL-schema that is specified explicitly or implicitly in the <object name> simply contained in <privileges> contained in the <revoke statement>.
  - 3) *PD* identifies the action identified by the <action> simply contained in <privileges> contained in the <revoke statement>.
  - 4) *PD* indicates that the privilege is grantable.

34) Specifications for Feature F052, “Intervals and datetime arithmetic”:

- a) Subclause 5.3, “<literal>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval literal>.
- b) Subclause 6.1, “<data type>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval type>.
- c) Subclause 6.27, “<numeric value function>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <extract expression>.
- d) Subclause 6.30, “<datetime value expression>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain <datetime value expression> that immediately contains a <plus sign> or a <minus sign>.
- e) Subclause 6.32, “<interval value expression>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval value expression>.
- f) Subclause 6.33, “<interval value function>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL shall not contain an <interval value function>.
- g) Subclause 10.1, “<interval qualifier>”:
  - i) Without Feature F052, “Intervals and datetime arithmetic”, conforming SQL language shall not contain an <interval qualifier>.

35) Specifications for Feature F053, “OVERLAPS predicate”:



a) Subclause 8.14, “<overlaps predicate>”:

- i) Without Feature F053, “OVERLAPS predicate”, conforming SQL language shall not contain an <overlaps predicate>.

36) Specifications for Feature F111, “Isolation levels other than SERIALIZABLE”:

a) Subclause 17.3, “<transaction characteristics>”:

- i) Without Feature F111, “Isolation levels other than SERIALIZABLE”, conforming SQL language shall not contain an <isolation level> that contains a <level of isolation> other than SERIALIZABLE.

37) Specifications for Feature F121, “Basic diagnostics management”:

a) Subclause 17.3, “<transaction characteristics>”:

- i) Without Feature F121, “Basic diagnostics management”, conforming SQL language shall not contain a <diagnostics size>.

b) Subclause 23.1, “<get diagnostics statement>”:

- i) Without Feature F121, “Basic diagnostics management”, conforming SQL language shall not contain a <get diagnostics statement>.

38) Specifications for Feature F122, “Enhanced diagnostics management”:

a) Subclause 23.1, “<get diagnostics statement>”:

- i) Without Feature F122, “Enhanced diagnostics management”, conforming SQL language shall not contain a <get diagnostics statement> containing a <simple target specification> whose declared type is different from that of the corresponding <statement information item name> or <condition information item name>.

39) Specifications for Feature F123, “All diagnostics”:

a) Subclause 23.1, “<get diagnostics statement>”:

- i) Without Feature F123, “All diagnostics”, conforming SQL language shall not contain an <all information>.

40) Specifications for Feature F171, “Multiple schemas per user”:

a) Subclause 11.1, “<schema definition>”:

- i) Without Feature F171, “Multiple schemas per user”, conforming SQL language shall not contain a <schema name clause> that contains a <schema name>.

41) Specifications for Feature F191, “Referential delete actions”:

a) Subclause 11.8, “<referential constraint definition>”:

- i) Without Feature F191, “Referential delete actions”, conforming SQL language shall not contain a <delete rule>.

42) Specifications for Feature F200, “TRUNCATE TABLE statement”:

a) Subclause 14.10, “<truncate table statement>”:

- i) Without Feature F200, “TRUNCATE TABLE statement”, conforming SQL language shall not contain a <truncate table statement>.

43) Specifications for Feature F222, “INSERT statement: DEFAULT VALUES clause”:

a) Subclause 14.11, “<insert statement>”:

- i) Without Feature F222, “INSERT statement: DEFAULT VALUES clause”, conforming SQL language shall not contain a <from default>.

44) Specifications for Feature F251, “Domain support”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <domain name>.

b) Subclause 6.4, “<value specification> and <target specification>”:

- i) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <general value specification> that contains VALUE.

c) Subclause 11.25, “<domain definition>”:

- i) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <domain definition>.

d) Subclause 11.31, “<drop domain statement>”:

- i) Without Feature F251, “Domain support”, conforming SQL language shall not contain a <drop domain statement>.

45) Specifications for Feature F262, “Extended CASE expression”:

a) Subclause 6.11, “<case expression>”:

- i) Without Feature F262, “Extended CASE expression”, in conforming SQL language, a <case operand> immediately contained in a <simple case> shall be a <row value predicand> that is a <row value constructor predicand> that is a single <common value expression> or <boolean predicand>.
- ii) Without Feature F262, “Extended CASE expression”, in conforming SQL language, a <when operand> contained in a <simple when clause> shall be a <row value predicand> that is a <row value constructor predicand> that is a single <common value expression> or <boolean predicand>.

46) Specifications for Feature F263, “Comma-separated predicates in simple CASE expression”:

a) Subclause 6.11, “<case expression>”:

- i) Without Feature F263, “Comma-separated predicates in simple CASE expression”, in conforming SQL language, a <when operand list> contained in a <simple when clause> shall simply contain exactly one <when operand>.

47) Specifications for Feature F271, “Compound character literals”:

a) Subclause 5.3, “<literal>”:

- i) Without Feature F271, “Compound character literals”, in conforming SQL language, a <character string literal> shall contain exactly one repetition of <character representation> (that is, it shall contain exactly one sequence of “<quote> <character representation>... <quote>”).

48) Specifications for Feature F281, “LIKE enhancements”:

a) Subclause 8.5, “<like predicate>”:

- i) Without Feature F281, “LIKE enhancements”, conforming SQL language shall not contain a <common value expression> simply contained in the <row value predicand> immediately contained in <character like predicate> that is not a column reference.
- ii) Without Feature F281, “LIKE enhancements”, conforming SQL language shall not contain a <character pattern> that is not a <value specification>.
- iii) Without Feature F281, “LIKE enhancements”, conforming SQL language shall not contain an <escape character> that is not a <value specification>.

49) Specifications for Feature F291, “UNIQUE predicate”:

a) Subclause 8.11, “<unique predicate>”:

- i) Without Feature F291, “UNIQUE predicate”, conforming SQL language shall not contain a <unique predicate>.

NOTE 482 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

50) Specifications for Feature F301, “CORRESPONDING in query expressions”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F301, “CORRESPONDING in query expressions”, conforming SQL language shall not contain a <query expression> that contains CORRESPONDING.

51) Specifications for Feature F302, “INTERSECT table operator”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F302, “INTERSECT table operator”, conforming SQL language shall not contain a <query term> that contains INTERSECT.

52) Specifications for Feature F304, “EXCEPT ALL table operator”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F304, “EXCEPT ALL table operator”, conforming SQL language shall not contain a <query expression> that contains EXCEPT ALL.

NOTE 483 — If DISTINCT, INTERSECT or EXCEPT is specified, then the Conformance Rules of Subclause 9.10, “Grouping operations”, apply.

53) Specifications for Feature F312, “MERGE statement”:

a) Subclause 14.12, “<merge statement>”:

- i) Without Feature F312, “MERGE statement”, conforming SQL language shall not contain a <merge statement>.

54) Specifications for Feature F313, “Enhanced MERGE statement”:

a) Subclause 14.12, “<merge statement>”:

- i) Without Feature F313, “Enhanced MERGE statement”, in conforming SQL language, a <merge statement> shall not contain each of <merge when matched clause> and <merge when not matched clause> more than once.
- ii) Without Feature F313, “Enhanced MERGE statement”, in conforming SQL language, a <merge when matched clause> or a <merge when not matched clause> shall not immediately contain a <search condition>.

55) Specifications for Feature F321, “User authorization”:

a) Subclause 6.4, “<value specification> and <target specification>”:

- i) Without Feature F321, “User authorization”, conforming SQL language shall not contain a <general value specification> that contains CURRENT\_USER, SYSTEM\_USER, or SESSION\_USER.

NOTE 484 — Although CURRENT\_USER and USER are semantically the same, without Feature F321, “User authorization”, CURRENT\_USER shall be specified as USER.

b) Subclause 11.5, “<default clause>”:

- i) Without Feature F321, “User authorization”, conforming SQL language shall not contain a <default option> that contains CURRENT\_USER, SESSION\_USER, or SYSTEM\_USER.

NOTE 485 — Although CURRENT\_USER and USER are semantically the same, without Feature F321, “User authorization”, CURRENT\_USER shall be specified as USER.

c) Subclause 19.2, “<set session user identifier statement>”:

- i) Without Feature F321, “User authorization”, conforming SQL language shall not contain a <set session user identifier statement>.

56) Specifications for Feature F361, “Subprogram support”:

a) Subclause 21.1, “<embedded SQL host program>”:

- i) Without Feature F361, “Subprogram support”, conforming SQL language shall not contain two <host variable definition>s that specify the same variable name.

57) Specifications for Feature F381, “Extended schema manipulation”:

a) Subclause 11.2, “<drop schema statement>”:

- i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop schema statement>.

b) Subclause 11.12, “<alter column definition>”:

- i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <alter column definition>.

c) Subclause 11.13, “<set column default clause>”:

- i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <set column default clause>.

d) Subclause 11.14, “<drop column default clause>”:

- i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop column default clause>.
- e) Subclause 11.15, “<add column scope clause>”:
  - i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <add column scope clause>.
- f) Subclause 11.16, “<drop column scope clause>”:
  - i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop column scope clause>.
- g) Subclause 11.20, “<add table constraint definition>”:
  - i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <add table constraint definition>.
- h) Subclause 11.21, “<drop table constraint definition>”:
  - i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain a <drop table constraint definition>.
- i) Subclause 11.52, “<alter routine statement>”:
  - i) Without Feature F381, “Extended schema manipulation”, conforming SQL language shall not contain an <alter routine statement>.

58) Specifications for Feature F382, “Alter column data type”:

- a) Subclause 11.17, “<alter column data type clause>”:
  - i) Without Feature F382, “Alter column data type”, conforming SQL language shall not contain an <alter column data type clause>.

59) Specifications for Feature F391, “Long identifiers”:

- a) Subclause 5.2, “<token> and <separator>”:
  - i) Without Feature F391, “Long identifiers”, in a <regular identifier>, the number of <identifier part>s shall be less than 18.
  - ii) Without Feature F391, “Long identifiers”, the <delimited identifier body> of a <delimited identifier> shall not comprise more than 18 <delimited identifier part>s.

NOTE 486 — Not every character set supported by a conforming SQL-implementation necessarily contains every character associated with <identifier start> and <identifier part> that is identified in the Syntax Rules of this Subclause. No conforming SQL-implementation shall be required to support in <identifier start> or <identifier part> any character identified in the Syntax Rules of this Subclause unless that character belongs to the character set in use for an SQL-client module or in SQL-data.

60) Specifications for Feature F392, “Unicode escapes in identifiers”:

- a) Subclause 5.2, “<token> and <separator>”:
  - i) Without Feature F392, “Unicode escapes in identifiers”, conforming SQL language shall not contain a <Unicode delimited identifier>.

61) Specifications for Feature F393, “Unicode escapes in literals”:

a) [Subclause 5.3](#), “<literal>”:

- i) Without Feature F393, “Unicode escapes in literals”, conforming SQL language shall not contain a <Unicode character string literal>.

62) Specifications for Feature F394, “Optional normal form specification”:

a) [Subclause 6.29](#), “<string value function>”:

- i) Without Feature F394, “Optional normal form specification”, conforming SQL language shall not contain <normal form>.

b) [Subclause 8.12](#), “<normalized predicate>”:

- i) Without Feature F394, “Optional normal form specification”, conforming SQL language shall not contain <normal form>.

63) Specifications for Feature F401, “Extended joined table”:

a) [Subclause 7.7](#), “<joined table>”:

- i) Without Feature F401, “Extended joined table”, conforming SQL language shall not contain a <cross join>.
- ii) Without Feature F401, “Extended joined table”, conforming SQL language shall not contain a <natural join>.
- iii) Without Feature F401, “Extended joined table”, conforming SQL language shall not contain FULL.

64) Specifications for Feature F402, “Named column joins for LOBs, arrays, and multisets”:

a) [Subclause 7.7](#), “<joined table>”:

- i) Without Feature F402, “Named column joins for LOBs, arrays, and multisets”, conforming SQL language shall not contain a <joined table> that simply contains either <natural join> or <named columns join> in which, if *C* is a corresponding join column, the declared type of *C* is LOB-ordered, array-ordered, or multiset-ordered.

NOTE 487 — If *C* is a corresponding join column, then the Conformance Rules of [Subclause 9.9](#), “Equality operations”, also apply.

65) Specifications for Feature F403, “Partitioned join tables”:

a) [Subclause 7.7](#), “<joined table>”:

- i) Without Feature F403, “Partitioned join tables”, conforming SQL language shall not contain <partitioned join table>.

66) Specifications for Feature F411, “Time zone specification”:

a) [Subclause 5.3](#), “<literal>”:

- i) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <time zone interval>.

b) [Subclause 6.1](#), “<data type>”:

- i) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain <with or without time zone>.
- c) Subclause 6.27, “<numeric value function>”:
  - i) Feature F411, “Time zone specification”, conforming SQL language shall not contain an <extract expression> that specifies a <time zone field>.
- d) Subclause 6.30, “<datetime value expression>”:
  - i) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <time zone>.
- e) Subclause 6.31, “<datetime value function>”:
  - i) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <current time value function>.
  - ii) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <current timestamp value function>.
- f) Subclause 19.4, “<set local time zone statement>”:
  - i) Without Feature F411, “Time zone specification”, conforming SQL language shall not contain a <set local time zone statement>.

67) Specifications for Feature F421, “National character”:

- a) Subclause 5.3, “<literal>”:
  - i) Without Feature F421, “National character”, conforming SQL language shall not contain a <national character string literal>.
- b) Subclause 6.1, “<data type>”:
  - i) Without Feature F421, “National character”, conforming SQL language shall not contain a <national character string type>
- c) Subclause 6.12, “<cast specification>”:
  - i) Without Feature F421, “National character”, conforming SQL language shall not contain a <cast operand> whose declared type is NATIONAL CHARACTER LARGE OBJECT.
- d) Subclause 6.27, “<numeric value function>”:
  - i) Without Feature F421, “National character”, conforming SQL language shall not contain a <length expression> that simply contains a <string value expression> that has a declared type of NATIONAL CHARACTER LARGE OBJECT.
- e) Subclause 8.5, “<like predicate>”:
  - i) Without Feature F421, “National character”, and Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <like predicate> shall not be of declared type NATIONAL CHARACTER LARGE OBJECT.

68) Specifications for Feature F431, “Read-only scrollable cursors”:

- a) Subclause 14.2, “<cursor properties>”:



- i) Without Feature F431, “Read-only scrollable cursors”, conforming SQL language shall not contain a <cursor scrollability>.
- b) Subclause 14.5, “<fetch statement>”:
  - i) Without Feature F431, “Read-only scrollable cursors”, a <fetch statement> shall not contain a <fetch orientation>.

69) Specifications for Feature F441, “Extended set function support”:

- a) Subclause 7.8, “<where clause>”:
  - i) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <value expression> directly contained in a <where clause> that contains a <column reference> that references a <derived column> that generally contains a <set function specification> without an intervening <routine invocation>.
- b) Subclause 10.9, “<aggregate function>”:
  - i) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <general set function> that contains a <computational operation> that immediately contains COUNT and does not contain a <set quantifier> that immediately contains DISTINCT.
  - ii) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <general set function> that does not contain a <set quantifier> that immediately contains DISTINCT and that contains a <value expression> that contains a column reference that does not reference a column of *T*.
  - iii) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <binary set function> that does not contain either a <dependent variable expression> or an <independent variable expression> that contains a column reference that references a column of *T*.
  - iv) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <value expression> simply contained in a <general set function> that contains a column reference that is an outer reference where the <value expression> is not a column reference.
  - v) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a <numeric value expression> simply contained in a <dependent variable expression> or an <independent variable expression> that contains a column reference that is an outer reference and in which the <numeric value expression> is not a column reference.
  - vi) Without Feature F441, “Extended set function support”, conforming SQL language shall not contain a column reference contained in an <aggregate function> that contains a reference to a column derived from a <value expression> that generally contains an <aggregate function> *SFS2* without an intervening <routine invocation>.

70) Specifications for Feature F442, “Mixed column references in set functions”:

- a) Subclause 10.9, “<aggregate function>”:
  - i) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain a <hypothetical set function value expression list> or a <sort specification list> that simply contains a <value expression> that contains more than one column reference, one of which is an outer reference.



- ii) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain an <inverse distribution function> that contains an <inverse distribution function argument> or a <sort specification> that contains more than one column reference, one of which is an outer reference.
- iii) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain an <aggregate function> that contains a <general set function> whose simply contained <value expression> contains more than one column reference, one of which is an outer reference.
- iv) Without Feature F442, “Mixed column references in set functions”, conforming SQL language shall not contain an <aggregate function> that contains a <binary set function> whose simply contained <dependent variable expression> or <independent variable expression> contains more than one column reference, one of which is an outer reference.

71) Specifications for Feature F451, “Character set definition”:

- a) Subclause 11.32, “<character set definition>”:
  - i) Without Feature F451, “Character set definition”, conforming SQL language shall not contain a <character set definition>.
- b) Subclause 11.33, “<drop character set statement>”:
  - i) Without Feature F451, “Character set definition”, conforming SQL language shall not contain a <drop character set statement>.

72) Specifications for Feature F461, “Named character sets”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <character set name>.
- b) Subclause 10.5, “<character set specification>”:
  - i) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <character set specification>.
- c) Subclause 11.1, “<schema definition>”:
  - i) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <schema character set specification>.
- d) Subclause 13.2, “<module name clause>”:
  - i) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <module character set specification>.
- e) Subclause 19.7, “<set names statement>”:
  - i) Without Feature F461, “Named character sets”, conforming SQL language shall not contain a <set names statement>.
- f) Subclause 21.1, “<embedded SQL host program>”:
  - i) Without Feature F461, “Named character sets”, conforming SQL language shall not contain an <embedded character set declaration>.

73) Specifications for Feature F491, “Constraint management”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature F491, “Constraint management”, conforming SQL language shall not contain a <constraint name>.
- b) Subclause 10.8, “<constraint name definition> and <constraint characteristics>”:
  - i) Without Feature F491, “Constraint management”, conforming SQL language shall not contain a <constraint name definition>.
- c) Subclause 11.30, “<drop domain constraint definition>”:
  - i) Without Feature F491, “Constraint management”, conforming SQL language shall not contain a <drop domain constraint definition>.
- d) Subclause 21.2, “<embedded exception declaration>”:
  - i) Without Feature F491, “Constraint management”, conforming SQL language shall not contain an <SQL condition> that contains a <constraint name>.

74) Specifications for Feature F521, “Assertions”:

- a) Subclause 11.38, “<assertion definition>”:
  - i) Without Feature F521, “Assertions”, conforming SQL language shall not contain an <assertion definition>.
- b) Subclause 11.39, “<drop assertion statement>”:
  - i) Without Feature F521, “Assertions”, conforming SQL language shall not contain a <drop assertion statement>.

75) Specifications for Feature F531, “Temporary tables”:

- a) Subclause 11.3, “<table definition>”:
  - i) Without Feature F531, “Temporary tables”, conforming SQL language shall not contain a <table scope> and shall not reference any global or local temporary table.
- b) Subclause 14.16, “<temporary table declaration>”:
  - i) Without Feature F531, “Temporary tables”, conforming SQL language shall not contain a <temporary table declaration>.

76) Specifications for Feature F555, “Enhanced seconds precision”:

- a) Subclause 5.3, “<literal>”:
  - i) Without Feature F555, “Enhanced seconds precision”, in conforming SQL language, an <unsigned integer> that is a <seconds fraction> that is contained in a <timestamp literal> shall not contain more than 6 <digit>s.
  - ii) Without Feature F555, “Enhanced seconds precision”, in conforming SQL language, a <time literal> shall not contain a <seconds fraction>.
- b) Subclause 6.1, “<data type>”:

- i) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <time precision> that does not specify 0 (zero).
- ii) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <timestamp precision> that does not specify either 0 (zero) or 6.
- c) Subclause 6.31, “<datetime value function>”:
  - i) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <current local time value function> that contains a <time precision> that is not 0 (zero).
  - ii) Without Feature F555, “Enhanced seconds precision”, conforming SQL language shall not contain a <current local timestamp value function> that contains a <timestamp precision> that is neither 0 (zero) nor 6.

77) Specifications for Feature F561, “Full value expressions”:

- a) Subclause 8.4, “<in predicate>”:
  - i) Without Feature F561, “Full value expressions”, conforming SQL language shall not contain a <row value expression> immediately contained in an <in value list> that is not a <value specification>.
 

NOTE 488 — Since <in predicate> is an equality operation, the Conformance Rules of Subclause 9.9, “Equality operations”, also apply.
- b) Subclause 10.9, “<aggregate function>”:
  - i) Without Feature F561, “Full value expressions”, or Feature F801, “Full set function”, conforming SQL language shall not contain a <general set function> that immediately contains DISTINCT and contains a <value expression> that is not a column reference.

78) Specifications for Feature F571, “Truth value tests”:

- a) Subclause 6.34, “<boolean value expression>”:
  - i) Without Feature F571, “Truth value tests”, conforming SQL language shall not contain a <boolean test> that simply contains a <truth value>.

79) Specifications for Feature F591, “Derived tables”:

- a) Subclause 7.6, “<table reference>”:
  - i) Without Feature F591, “Derived tables”, conforming SQL language shall not contain a <derived table>.

80) Specifications for Feature F611, “Indicator data types”:

- a) Subclause 6.4, “<value specification> and <target specification>”:
  - i) Without Feature F611, “Indicator data types”, in conforming SQL language, the declared types of <indicator parameter>s and <indicator variable>s shall be the same implementation-defined data type.

81) Specifications for Feature F641, “Row and table constructors”:

- a) Subclause 7.1, “<row value constructor>”:

- i) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain an <explicit row value constructor> that is not simply contained in a <table value constructor> and that contains more than one <row value constructor element>.
  - ii) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain an <explicit row value constructor> that is a <row subquery>.
  - iii) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain a <contextually typed row value constructor> that is not simply contained in a <contextually typed table value constructor> and that contains more than one <row value constructor element>.
  - iv) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain a <contextually typed row value constructor> that is a <row subquery>.
- b) Subclause 7.3, “<table value constructor>”:
- i) Without Feature F641, “Row and table constructors”, in conforming SQL language, the <contextually typed row value expression list> of a <contextually typed table value constructor> shall contain exactly one <contextually typed row value constructor> *RVE*. *RVE* shall be of the form “(<contextually typed row value constructor element list>)”.
  - ii) Without Feature F641, “Row and table constructors”, conforming SQL language shall not contain a <table value constructor>.

82) Specifications for Feature F651, “Catalog name qualifiers”:

- a) Subclause 5.4, “Names and identifiers”:
- i) Without Feature F651, “Catalog name qualifiers”, conforming SQL language shall not contain a <catalog name>.
- b) Subclause 19.5, “<set catalog statement>”:
- i) Without Feature F651, “Catalog name qualifiers”, conforming SQL language shall not contain a <set catalog statement>.

83) Specifications for Feature F661, “Simple tables”:

- a) Subclause 7.13, “<query expression>”:
- i) Without Feature F661, “Simple tables”, conforming SQL language shall not contain a <simple table> that immediately contains a <table value constructor> except in an <insert statement>.
  - ii) Without Feature F661, “Simple tables”, conforming SQL language shall not contain an <explicit table>.

84) Specifications for Feature F671, “Subqueries in CHECK constraints”:

- a) Subclause 11.9, “<check constraint definition>”:
- i) Without Feature F671, “Subqueries in CHECK constraints”, conforming SQL language shall not contain a <search condition> contained in a <check constraint definition> that contains a <query expression>.

85) Specifications for Feature F672, “Retrospective check constraints”:

- a) Subclause 11.9, “<check constraint definition>”:

- i) Without Feature F672, “Retrospective check constraints”, conforming SQL language shall not contain a <check constraint definition> that generally contains CURRENT\_DATE, CURRENT\_TIMESTAMP, or LOCALTIMESTAMP.
- b) Subclause 11.38, “<assertion definition>”:
  - i) Without Feature F672, “Retrospective check constraints”, conforming SQL language shall not contain an <assertion definition> that generally contains CURRENT\_DATE, CURRENT\_TIMESTAMP, or LOCALTIMESTAMP.

86) Specifications for Feature F690, “Collation support”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <collation name>.
- b) Subclause 10.7, “<collate clause>”:
  - i) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <collate clause>.
- c) Subclause 11.34, “<collation definition>”:
  - i) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <collation definition>.
- d) Subclause 11.35, “<drop collation statement>”:
  - i) Without Feature F690, “Collation support”, conforming SQL language shall not contain a <drop collation statement>.

87) Specifications for Feature F692, “Extended collation support”:

- a) Subclause 11.4, “<column definition>”:
  - i) Without Feature F692, “Extended collation support”, conforming SQL language shall not contain a <column definition> that immediately contains a <collate clause>.
- b) Subclause 11.25, “<domain definition>”:
  - i) Without Feature F692, “Extended collation support”, conforming SQL language shall not contain a <domain definition> that immediately contains a <collate clause>.
- c) Subclause 11.43, “<attribute definition>”:
  - i) Without Feature F692, “Extended collation support”, conforming SQL language shall not contain an <attribute definition> that immediately contains a <collate clause>.

88) Specifications for Feature F693, “SQL-session and client module collations”:

- a) Subclause 6.4, “<value specification> and <target specification>”:
  - i) Without Feature F693, “SQL-session and client module collations”, conforming SQL language shall not contain <current collation specification>.
- b) Subclause 13.1, “<SQL-client module definition>”:

- i) Without Feature F693, “SQL-session and client module collations”, conforming SQL language shall not contain a <module collation specification>.
- c) Subclause 19.10, “<set session collation statement>”:
  - i) Without Feature F693, “SQL-session and client module collations”, conforming SQL language shall not contain a <set session collation statement>.

89) Specifications for Feature F695, “Translation support”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transliteration name>.
  - ii) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transcoding name>.
- b) Subclause 6.29, “<string value function>”:
  - i) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <character transliteration>.
  - ii) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transcoding>.
- c) Subclause 11.36, “<transliteration definition>”:
  - i) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <transliteration definition>.
- d) Subclause 11.37, “<drop transliteration statement>”:
  - i) Without Feature F695, “Translation support”, conforming SQL language shall not contain a <drop transliteration statement>.

90) Specifications for Feature F701, “Referential update actions”:

- a) Subclause 11.8, “<referential constraint definition>”:
  - i) Without Feature F701, “Referential update actions”, conforming SQL language shall not contain an <update rule>.

91) Specifications for Feature F711, “ALTER domain”:

- a) Subclause 11.26, “<alter domain statement>”:
  - i) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain an <alter domain statement>.
- b) Subclause 11.27, “<set domain default clause>”:
  - i) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain a <set domain default clause>.
- c) Subclause 11.28, “<drop domain default clause>”:
  - i) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain a <drop domain default clause>.

d) Subclause 11.29, “<add domain constraint definition>”:

- i) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain an <add domain constraint definition>.

e) Subclause 11.30, “<drop domain constraint definition>”:

- i) Without Feature F711, “ALTER domain”, conforming SQL language shall not contain a <drop domain constraint definition>.

92) Specifications for Feature F721, “Deferrable constraints”:

a) Subclause 10.8, “<constraint name definition> and <constraint characteristics>”:

- i) Without Feature F721, “Deferrable constraints”, conforming SQL language shall not contain a <constraint characteristics>.

NOTE 489 — This means that INITIALLY IMMEDIATE NOT DEFERRABLE is implicit.

b) Subclause 17.4, “<set constraints mode statement>”:

- i) Without Feature F721, “Deferrable constraints”, conforming SQL language shall not contain a <set constraints mode statement>.

93) Specifications for Feature F731, “INSERT column privileges”:

a) Subclause 12.3, “<privileges>”:

- i) Without Feature F731, “INSERT column privileges”, in conforming SQL language, an <action> that contains INSERT shall not contain a <privilege column list>.

94) Specifications for Feature F741, “Referential MATCH types”:

a) Subclause 8.13, “<match predicate>”:

- i) Without Feature F741, “Referential MATCH types”, conforming SQL language shall not contain a <match predicate>.

NOTE 490 — The Conformance Rules of Subclause 9.9, “Equality operations”, also apply.

b) Subclause 11.8, “<referential constraint definition>”:

- i) Without Feature F741, “Referential MATCH types”, conforming SQL language shall not contain a <references specification> that contains MATCH.

95) Specifications for Feature F751, “View CHECK enhancements”:

a) Subclause 11.23, “<view definition>”:

- i) Without Feature F751, “View CHECK enhancements”, conforming SQL language shall not contain a <levels clause>.
- ii) Without Feature F751, “View CHECK enhancements”, conforming SQL language shall not contain <view definition> that contains a <query expression> that contains a <query expression> and contains WITH CHECK OPTION.

96) Specifications for Feature F761, “Session management”:

a) Subclause 19.1, “<set session characteristics statement>”:



- i) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set session characteristics statement>.
- b) Subclause 19.5, “<set catalog statement>”:
  - i) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set catalog statement>.
- c) Subclause 19.6, “<set schema statement>”:
  - i) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set schema statement>.
- d) Subclause 19.7, “<set names statement>”:
  - i) Without Feature F761, “Session management”, conforming SQL language shall not contain a <set names statement>.

97) Specifications for Feature F762, “CURRENT\_CATALOG”:

- a) Subclause 6.4, “<value specification> and <target specification>”:
  - i) Without Feature F762, “CURRENT\_CATALOG”, conforming SQL language shall not contain a <general value specification> that contains CURRENT\_CATALOG.
- b) Subclause 11.5, “<default clause>”:
  - i) Without Feature F762, “CURRENT\_CATALOG”, conforming SQL language shall not contain a <default option> that contains CURRENT\_CATALOG.

98) Specifications for Feature F763, “CURRENT\_SCHEMA”:

- a) Subclause 6.4, “<value specification> and <target specification>”:
  - i) Without Feature F763, “CURRENT\_SCHEMA”, conforming SQL language shall not contain a <general value specification> that contains CURRENT\_SCHEMA.
- b) Subclause 11.5, “<default clause>”:
  - i) Without Feature F763, “CURRENT\_SCHEMA”, conforming SQL language shall not contain a <default option> that contains CURRENT\_SCHEMA.

99) Specifications for Feature F771, “Connection management”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature F771, “Connection management”, conforming SQL language shall not contain an explicit <connection name>.
- b) Subclause 18.1, “<connect statement>”:
  - i) Without Feature F771, “Connection management”, conforming SQL language shall not contain a <connect statement>.
- c) Subclause 18.2, “<set connection statement>”:
  - i) Without Feature F771, “Connection management”, conforming SQL language shall not contain a <set connection statement>.



d) Subclause 18.3, “<disconnect statement>”:

- i) Without Feature F771, “Connection management”, conforming SQL language shall not contain a <disconnect statement>.

100) Specifications for Feature F781, “Self-referencing operations”:

a) Subclause 14.9, “<delete statement: searched>”:

- i) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <delete statement: searched> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in the <search condition>.

b) Subclause 14.11, “<insert statement>”:

- i) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain an <insert statement> in which the <table name> of a leaf generally underlying table of *T* is generally contained in the <from subquery> except as the table name of a qualifying table of a column reference.

c) Subclause 14.12, “<merge statement>”:

- i) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <merge statement> in which a leaf generally underlying table of *T* is generally contained in a <query expression> immediately contained in the <table reference> except as the <table or query name> or <correlation name> of a column reference.
- ii) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <merge statement> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in the <search condition>.

d) Subclause 14.14, “<update statement: searched>”:

- i) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain an <update statement: searched> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in the <search condition>.

e) Subclause 14.15, “<set clause list>”:

- i) Without Feature F781, “Self-referencing operations”, conforming SQL language shall not contain a <set clause> in which a leaf generally underlying table of *T* is an underlying table of any <query expression> generally contained in any <value expression> simply contained in an <update source> or <assigned row> immediately contained in the <set clause>.

101) Specifications for Feature F791, “Insensitive cursors”:

a) Subclause 14.2, “<cursor properties>”:

- i) Without Feature F791, “Insensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains INSENSITIVE.
- ii) Without Feature F791, “Insensitive cursors”, or Feature T231, “Sensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains ASENSITIVE.

102) Specifications for Feature F801, “Full set function”:

a) Subclause 7.12, “<query specification>”:

- i) Without Feature F801, “Full set function”, conforming SQL language shall not contain a <query specification> *QS* that contains more than one <set quantifier> containing DISTINCT, unless such <set quantifier> is contained in a <query expression> contained in *QS*.

103) Specifications for Feature F821, “Local table references”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature F821, “Local table references”, conforming SQL language shall not contain a <local or schema qualifier> that contains a <local qualifier>.

b) Subclause 6.7, “<column reference>”:

- i) Without Feature F821, “Local table references”, conforming SQL language shall not contain a <column reference> that simply contains MODULE.

104) Specifications for Feature F831, “Full cursor update”:

a) Subclause 14.1, “<declare cursor>”:

- i) Without Feature F831, “Full cursor update”, conforming SQL language shall not contain a <declare cursor> that contains both a <cursor specification> that contains an <updatability clause> that specifies FOR UPDATE and <cursor properties> that contain a <cursor scrollability>.

b) Subclause 14.3, “<cursor specification>”:

- i) Without Feature F831, “Full cursor update”, conforming SQL language shall not contain a <cursor specification> that contains both an <updatability clause> that specifies FOR UPDATE and an <order by clause>.

c) Subclause 14.13, “<update statement: positioned>”:

- i) Without Feature F831, “Full cursor update”, conforming SQL language shall not contain an <update statement: positioned> in which the declared <cursor specification> of *CDD* is ordered.

105) Specifications for Feature F841, “LIKE\_REGEX predicate”:

a) Subclause 8.7, “<regex like predicate>”:

- i) Without Feature F841, “LIKE\_REGEX predicate”, conforming SQL language shall not contain <regex like predicate>.

106) Specifications for Feature F842, “OCCURENCES\_REGEX function”:

a) Subclause 6.27, “<numeric value function>”:

- i) Without Feature F842, “OCCURENCES\_REGEX function”, conforming SQL language shall not contain <regex occurrences function>.

107) Specifications for Feature F843, “POSITION\_REGEX function”:

a) Subclause 6.27, “<numeric value function>”:

- i) Without Feature F843, “POSITION\_REGEX function”, conforming SQL language shall not contain <regex position expression>.

108) Specifications for Feature F844, “SUBSTRING\_REGEX”:

a) Subclause 6.29, “<string value function>”:

- i) Without Feature F844, “SUBSTRING\_REGEX”, conforming SQL language shall not contain <regex substring function>.

109) Specifications for Feature F845, “TRANSLATE\_REGEX”:

- a) Subclause 6.29, “<string value function>”:
  - i) Without Feature F845, “TRANSLATE\_REGEX”, conforming SQL language shall not contain <regex transliteration>.

110) Specifications for Feature F846, “Octet support in regular expression operators”:

- a) Subclause 6.27, “<numeric value function>”:
  - i) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex occurrences function> shall not contain <char length units>.
  - ii) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex position expression> shall not contain <char length units>.
- b) Subclause 6.29, “<string value function>”:
  - i) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex substring function> shall not contain <char length units>.
  - ii) Without Feature F846, “Octet support in regular expression operators”, in conforming SQL language, <regex transliteration> shall not contain <char length units>.

111) Specifications for Feature F847, “Nonconstant regular expression”:

- a) Subclause 6.27, “<numeric value function>”:
  - i) Without Feature F847, “Nonconstant regular expression”, in conforming SQL language, <XQuery pattern> and <XQuery option flag> shall be <value specification>s.
- b) Subclause 6.29, “<string value function>”:
  - i) Without Feature F847, “Nonconstant regular expressions”, in conforming SQL language, <XQuery pattern>, <XQuery option flag>, and <XQuery replacement string> shall be <value specification>s.
- c) Subclause 8.7, “<regex like predicate>”:
  - i) Without Feature F847, “Nonconstant regular expressions”, in conforming SQL language, <XQuery pattern> and <XQuery option flag> shall be <value specification>s.

112) Specifications for Feature F850, “Top-level <order by clause> in <query expression>”:

- a) Subclause 7.13, “<query expression>”:
  - i) Without Feature F850, “Top-level <order by clause> in <query expression>”, in conforming SQL language, a <query expression> not immediately contained in either an <array value constructor by query> or a <cursor specification> shall not immediately contain an <order by clause>.

113) Specifications for Feature F851, “<order by clause> in subqueries”:

- a) Subclause 7.13, “<query expression>”:

- i) Without Feature F851, “<order by clause> in subqueries”, in conforming SQL language, a <query expression> contained in another <query expression> shall not immediately contain an <order by clause>.

114) Specifications for Feature F852, “Top-level <order by clause> in views”:

a) Subclause 11.23, “<view definition>”:

- i) Without Feature F852, “Top-level <order by clause> in views”, in conforming SQL language, a <query expression> immediately contained in a <view definition> shall not immediately contain an <order by clause>.

115) Specifications for Feature F855, “Nested <order by clause> in <query expression>”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F855, “Nested <order by clause> in <query expression>”, in conforming SQL language, a <query primary> shall not immediately contain an <order by clause>.

116) Specifications for Feature F856, “Nested <fetch first clause> in <query expression>”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F856, “Nested <fetch first clause> in <query expression>”, in conforming SQL language, a <query primary> shall not immediately contain a <fetch first clause>.

117) Specifications for Feature F857, “Top-level <fetch first clause> in <query expression>”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F857, “Top-level <fetch first clause> in <query expression>”, in conforming SQL language, a <query expression> shall not immediately contain a <fetch first clause>.

118) Specifications for Feature F858, “<fetch first clause> in subqueries”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature F858, “<fetch first clause> in subqueries”, in conforming SQL language, a <query expression> contained in another <query expression> shall not immediately contain a <fetch first clause>.

119) Specifications for Feature F859, “Top-level <fetch first clause> in views”:

a) Subclause 11.23, “<view definition>”:

- i) Without Feature F859, “Top-level <fetch first clause> in views”, in conforming SQL language, a <query expression> immediately contained in a <view definition> shall not immediately contain a <fetch first clause>.

120) Specifications for Feature S023, “Basic structured types”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <attribute name>.

b) Subclause 6.1, “<data type>”:

- i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <path-resolved user-defined type name> that identifies a structured type.
- c) Subclause 6.16, “<method invocation>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <method invocation>.
- d) Subclause 6.18, “<new specification>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <new specification>.
- e) Subclause 10.4, “<routine invocation>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <generalized expression>.
- f) Subclause 11.42, “<user-defined type definition>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <member list>.
  - ii) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <method specification list>.
- g) Subclause 11.43, “<attribute definition>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain an <attribute definition>.
- h) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <method specification designator>.
- i) Subclause 12.3, “<privileges>”:
  - i) Without Feature S023, “Basic structured types”, conforming SQL language shall not contain a <privileges> that contains an <action> that contains UNDER and that contains an <object name> that contains a <schema-resolved user-defined type name> that identifies a structured type.

12l) Specifications for Feature S024, “Enhanced structured types”:

- a) Subclause 6.17, “<static method invocation>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <static method invocation>.
- b) Subclause 9.9, “Equality operations”:
  - i) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared type of an operand of an equality operation shall not be ST-ordered.
- c) Subclause 9.10, “Grouping operations”:
  - i) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared type of an operand of a grouping operation shall not be ST-ordered.

- d) Subclause 9.11, “Multiset element grouping operations”:
  - i) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared element type of a multiset operand of a multiset element grouping operation shall not be ST-ordered.
- e) Subclause 9.12, “Ordering operations”:
  - i) Without Feature S024, “Enhanced structured types”, in conforming SQL language, the declared type of an operand of an ordering operation shall not be ST-ordered.
- f) Subclause 10.6, “<specific routine designator>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <specific routine designator> that contains a <routine type> that immediately contains METHOD.
- g) Subclause 11.42, “<user-defined type definition>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <instantiable clause> that contains NOT INSTANTIABLE.
  - ii) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <original method specification> that immediately contains SELF AS RESULT.
  - iii) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <method characteristics> that contains a <parameter style> that contains GENERAL.
  - iv) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <original method specification> that contains an <SQL-data access indication> that immediately contains NO SQL.
  - v) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <partial method specification> that contains INSTANCE or STATIC.
- h) Subclause 11.43, “<attribute definition>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <attribute default>.
- i) Subclause 11.44, “<alter type statement>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain an <alter type statement>.
- j) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature S024, “Enhanced structured types”, an <SQL parameter declaration> shall not contain RESULT.
- k) Subclause 11.53, “<drop routine statement>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <drop routine statement> that contains a <specific routine designator> that identifies a method.
- l) Subclause 12.2, “<grant privilege statement>”:

- i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <specific routine designator> contained in a <grant privilege statement> that identifies a method.
- m) Subclause 12.3, “<privileges>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <privileges> that contains an <action> that contains USAGE and that contains an <object name> that contains a <schema-resolved user-defined type name> that identifies a structured type.
  - ii) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <privilege method list>.
- n) Subclause 14.11, “<insert statement>”:
  - i) Without Feature S024, “Enhanced structured types”, in conforming SQL language, for each column *C* identified in the explicit or implicit <insert column list>, if the declared type of *C* is a structured type *TY*, then the declared type of the corresponding column of the <query expression> or <contextually typed table value constructor> shall be *TY*.
- o) Subclause 14.12, “<merge statement>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <merge statement> that does not satisfy the condition: for each column *C* identified in the explicit or implicit <insert column list>, if the declared type of *C* is a structured type *TY*, then the declared type of the corresponding column of the <query expression> or <contextually typed table value constructor> is *TY*.
- p) Subclause 14.15, “<set clause list>”:
  - i) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <set clause> in which the declared type of the <update target> in the <set clause> is a structured type *TY* and the declared type of the <update source> or corresponding field of the <assigned row> contained in the <set clause> is not *TY*.
  - ii) Without Feature S024, “Enhanced structured types”, conforming SQL language shall not contain a <set clause> that contains a <mutated set clause> and in which the declared type of the last <method name> identifies a structured type *TY*, and the declared type of the <update source> contained in the <set clause> is not *TY*.

## 122) Specifications for Feature S025, “Final structured types”:

- a) Subclause 11.42, “<user-defined type definition>”:
  - i) Without Feature S025, “Final structured types”, in conforming SQL language, a <user-defined type definition> that defines a structured type shall contain a <finality> that is NOT FINAL.

## 123) Specifications for Feature S026, “Self-referencing structured types”:

- a) Subclause 11.43, “<attribute definition>”:
  - i) Without Feature S026, “Self-referencing structured types”, conforming SQL language shall not contain a <data type> simply contained in an <attribute definition> that is a <reference type> whose <referenced type> is equivalent to the <schema-resolved user-defined type name> simply contained in the <user-defined type definition> that contains <attribute definition>.



124) Specifications for Feature S027, “Create method by specific method name”:

- a) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature S027, “Create method by specific method name”, conforming SQL language shall not contain a <method specification designator> that contains SPECIFIC METHOD.

125) Specifications for Feature S028, “Permutable UDT options list”:

- a) Subclause 11.42, “<user-defined type definition>”:
  - i) Without Feature S028, “Permutable UDT options list”, conforming SQL language shall not contain a <user-defined type option list> in which <instantiable clause>, if specified, <finality>, <reference type specification>, if specified, <cast to ref>, if specified, <cast to type>, if specified, <cast to distinct>, if specified, and <cast to source>, if specified, do not appear in that sequence.

126) Specifications for Feature S041, “Basic reference types”:

- a) Subclause 6.1, “<data type>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <reference type>.
- b) Subclause 6.19, “<attribute or method reference>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain an <attribute or method reference>.
- c) Subclause 6.20, “<dereference operation>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <dereference operation>.
- d) Subclause 6.25, “<value expression>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <reference value expression>.
- e) Subclause 21.3, “<embedded SQL Ada program>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain an <Ada REF variable>.
- f) Subclause 21.4, “<embedded SQL C program>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <C REF variable>.
- g) Subclause 21.5, “<embedded SQL COBOL program>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <COBOL REF variable>.
- h) Subclause 21.6, “<embedded SQL Fortran program>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <Fortran REF variable>.
- i) Subclause 21.7, “<embedded SQL MUMPS program>”:



- i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <MUMPS REF variable>.
- j) Subclause 21.8, “<embedded SQL Pascal program>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <Pascal REF variable>.
- k) Subclause 21.9, “<embedded SQL PL/I program>”:
  - i) Without Feature S041, “Basic reference types”, conforming SQL language shall not contain a <PL/I REF variable>.

127) Specifications for Feature S043, “Enhanced reference types”:

- a) Subclause 6.1, “<data type>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <scope clause> that is not simply contained in a <data type> that is simply contained in a <column definition>.
- b) Subclause 6.12, “<cast specification>”:
  - i) Without Feature S043, “Enhanced reference types”, in conforming SQL language, if the declared type of <cast operand> is a reference type, then <cast target> shall contain a <data type> that is a reference type.
- c) Subclause 6.21, “<method reference>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <method reference>.
- d) Subclause 6.22, “<reference resolution>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <reference resolution>.
- e) Subclause 11.3, “<table definition>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <column option list> that contains a <scope clause>.
  - ii) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain <reference generation> that does not contain SYSTEM GENERATED.
- f) Subclause 11.15, “<add column scope clause>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain an <add column scope clause>.
- g) Subclause 11.16, “<drop column scope clause>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <drop column scope clause>.
- h) Subclause 11.23, “<view definition>”:
  - i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <referenceable view specification>.

i) Subclause 11.42, “<user-defined type definition>”:

- i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain a <reference type specification>.

j) Subclause 14.11, “<insert statement>”:

- i) Without Feature S043, “Enhanced reference types”, conforming SQL language shall not contain an <override clause>.

128) Specifications for Feature S051, “Create table of type”:

a) Subclause 11.3, “<table definition>”:

- i) Without Feature S051, “Create table of type”, conforming SQL language shall not contain “OF <path-resolved user-defined type name>”.

129) Specifications for Feature S071, “SQL paths in function and type name resolution”:

a) Subclause 6.4, “<value specification> and <target specification>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <general value specification> that contains CURRENT\_PATH.

b) Subclause 10.3, “<path specification>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <path specification>.

c) Subclause 11.1, “<schema definition>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <schema path specification>.

d) Subclause 11.5, “<default clause>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <default option> that contains CURRENT\_PATH.

e) Subclause 13.1, “<SQL-client module definition>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain a <module path specification>.

f) Subclause 19.8, “<set path statement>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, Conforming SQL language shall not contain a <set path statement>.

g) Subclause 21.1, “<embedded SQL host program>”:

- i) Without Feature S071, “SQL paths in function and type name resolution”, conforming SQL language shall not contain an <embedded path specification>.

130) Specifications for Feature S081, “Subtables”:

a) Subclause 11.3, “<table definition>”:

- i) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <subtable clause>.
- b) Subclause 12.2, “<grant privilege statement>”:
  - i) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <grant privilege statement> that contains WITH HIERARCHY OPTION.
- c) Subclause 12.3, “<privileges>”:
  - i) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <privileges> that contains an <action> that contains UNDER and that contains an <object name> that contains a <table name>.
- d) Subclause 12.7, “<revoke statement>”:
  - i) Without Feature S081, “Subtables”, conforming SQL language shall not contain a <revoke option extension> that contains HIERARCHY OPTION FOR.

13l) Specifications for Feature S091, “Basic array support”:

- a) Subclause 6.1, “<data type>”:
  - i) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array type>.
- b) Subclause 6.5, “<contextually typed value specification>”:
  - i) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <empty specification> that simply contains ARRAY.
- c) Subclause 6.23, “<array element reference>”:
  - i) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array element reference>.
- d) Subclause 6.27, “<numeric value function>”:
  - i) Without Feature S091, “Basic array support”, or Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <cardinality expression>.
- e) Subclause 6.35, “<array value expression>”:
  - i) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array value expression>.
- f) Subclause 6.36, “<array value constructor>”:
  - i) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <array value constructor by enumeration>.
- g) Subclause 7.6, “<table reference>”:
  - i) Without Feature S091, “Basic array support”, or Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <collection derived table>.
- h) Subclause 14.15, “<set clause list>”:

- i) Without Feature S091, “Basic array support”, conforming SQL language shall not contain an <update target> that immediately contains a <simple value specification>.

132) Specifications for Feature S092, “Arrays of user-defined types”:

a) Subclause 6.1, “<data type>”:

- i) Without Feature S092, “Arrays of user-defined types”, conforming SQL language shall not contain an <array type> that is based on a <data type> that contains a <path-resolved user-defined type name>.

133) Specifications for Feature S094, “Arrays of reference types”:

a) Subclause 6.1, “<data type>”:

- i) Without Feature S094, “Arrays of reference types”, conforming SQL language shall not contain an <array type> that is based on a <data type> that contains a <reference type>.

134) Specifications for Feature S095, “Array constructors by query”:

a) Subclause 6.36, “<array value constructor>”:

- i) Without Feature S095, “Array constructors by query”, conforming SQL language shall not contain an <array value constructor by query>.

135) Specifications for Feature S096, “Optional array bounds”:

a) Subclause 6.1, “<data type>”:

- i) Without Feature S096, “Optional array bounds”, conforming SQL language shall not contain an <array type> that does not immediately contain <maximum cardinality>.

136) Specifications for Feature S097, “Array element assignment”:

a) Subclause 6.4, “<value specification> and <target specification>”:

- i) Without Feature S097, “Array element assignment”, conforming SQL language shall not contain a <target array element specification>.

137) Specifications for Feature S098, “ARRAY\_AGG”:

a) Subclause 10.9, “<aggregate function>”:

- i) Without Feature S098, “ARRAY\_AGG”, conforming SQL language shall not contain an <array aggregate function>.

138) Specifications for Feature S111, “ONLY in query expressions”:

a) Subclause 7.6, “<table reference>”:

- i) Without Feature S111, “ONLY in query expressions”, conforming SQL language shall not contain a <table reference> that contains an <only spec>.

b) Subclause 14.8, “<delete statement: positioned>”:

- i) Without Feature S111, “ONLY in query expressions”, conforming SQL language shall not contain a <target table> that contains ONLY.

139) Specifications for Feature S151, “Type predicate”:

a) Subclause 8.19, “<type predicate>”:

- i) Without Feature S151, “Type predicate”, conforming SQL language shall not contain a <type predicate>.

140) Specifications for Feature S161, “Subtype treatment”:

a) Subclause 6.15, “<subtype treatment>”:

- i) Without Feature S161, “Subtype treatment”, conforming SQL Language shall not contain a <subtype treatment>.

141) Specifications for Feature S162, “Subtype treatment for references”:

a) Subclause 6.15, “<subtype treatment>”:

- i) Without Feature S162, “Subtype treatment for references”, conforming SQL language shall not contain a <target subtype> that contains a <reference type>.

142) Specifications for Feature S201, “SQL-invoked routines on arrays”:

a) Subclause 10.4, “<routine invocation>”:

- i) Without Feature S201, “SQL-invoked routines on arrays”, conforming SQL language shall not contain an <SQL argument> whose declared type is an array type.

b) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature S201, “SQL-invoked routines on arrays”, conforming SQL language shall not contain a <parameter type> that is based on an array type.
- ii) Without Feature S201, “SQL-invoked routines on arrays”, conforming SQL language shall not contain a <returns data type> that is based on an array type.

143) Specifications for Feature S202, “SQL-invoked routines on multisets”:

a) Subclause 10.4, “<routine invocation>”:

- i) Without Feature S202, “SQL-invoked routines on multisets”, conforming SQL language shall not contain an <SQL argument> whose declared type is a multiset type.

b) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature S202, “SQL-invoked routines on multisets”, conforming SQL language shall not contain a <parameter type> that is based on a multiset type.
- ii) Without Feature S202, “SQL-invoked routines on multisets”, conforming SQL language shall not contain a <returns data type> that is based on a multiset type.

144) Specifications for Feature S211, “User-defined cast functions”:

a) Subclause 11.54, “<user-defined cast definition>”:

- i) Without Feature S211, “User-defined cast functions”, conforming SQL language shall not contain a <user-defined cast definition>.

b) Subclause 11.55, “<drop user-defined cast statement>”:

- i) Without Feature S211, “User-defined cast functions”, conforming SQL language shall not contain a <drop user-defined cast statement>.

145) Specifications for Feature S231, “Structured type locators”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies a structured type.
- ii) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies a structured type.

b) Subclause 13.3, “<externally-invoked procedure>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <host parameter data type> that simply contains a <data type> that specifies a structured type and that contains a <locator indication>.

c) Subclause 21.3, “<embedded SQL Ada program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in an <Ada user-defined type locator variable> that identifies a structured type.

d) Subclause 21.4, “<embedded SQL C program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <C user-defined type locator variable> that identifies a structured type.

e) Subclause 21.5, “<embedded SQL COBOL program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <COBOL user-defined type locator variable> that identifies a structured type.

f) Subclause 21.6, “<embedded SQL Fortran program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <Fortran user-defined type locator variable> that identifies a structured type.

g) Subclause 21.7, “<embedded SQL MUMPS program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <MUMPS user-defined type locator variable> that identifies a structured type.

h) Subclause 21.8, “<embedded SQL Pascal program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <Pascal user-defined type locator variable> that identifies a structured type.

i) Subclause 21.9, “<embedded SQL PL/I program>”:

- i) Without Feature S231, “Structured type locators”, conforming SQL language shall not contain a <path-resolved user-defined type name> simply contained in a <PL/I user-defined type locator variable> that identifies a structured type.

146) Specifications for Feature S232, “Array locators”:

- a) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies an array type.
  - ii) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies an array type.
- b) Subclause 13.3, “<externally-invoked procedure>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <host parameter data type> that simply contains an <array type> and that contains a <locator indication>.
- c) Subclause 21.3, “<embedded SQL Ada program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain an <Ada array locator variable>.
- d) Subclause 21.4, “<embedded SQL C program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain an <C array locator variable>.
- e) Subclause 21.5, “<embedded SQL COBOL program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <COBOL array locator variable>.
- f) Subclause 21.6, “<embedded SQL Fortran program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <Fortran array locator variable>.
- g) Subclause 21.7, “<embedded SQL MUMPS program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <MUMPS array locator variable>.
- h) Subclause 21.8, “<embedded SQL Pascal program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <Pascal array locator variable>.
- i) Subclause 21.9, “<embedded SQL PL/I program>”:
  - i) Without Feature S232, “Array locators”, conforming SQL language shall not contain a <PL/I array locator variable>.

147) Specifications for Feature S233, “Multiset locators”:



- a) Subclause 11.51, “<SQL-invoked routine>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies a multiset type.
    - ii) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies a multiset type.
  - b) Subclause 13.3, “<externally-invoked procedure>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <host parameter data type> that simply contains a <multiset type> and that contains a <locator indication>.
  - c) Subclause 21.3, “<embedded SQL Ada program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain an <Ada multiset locator variable>.
  - d) Subclause 21.4, “<embedded SQL C program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <C multiset locator variable>.
  - e) Subclause 21.5, “<embedded SQL COBOL program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <COBOL multiset locator variable>.
  - f) Subclause 21.6, “<embedded SQL Fortran program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <Fortran multiset locator variable>.
  - g) Subclause 21.7, “<embedded SQL MUMPS program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <MUMPS multiset locator variable>.
  - h) Subclause 21.8, “<embedded SQL Pascal program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <Pascal multiset locator variable>.
  - i) Subclause 21.9, “<embedded SQL PL/I program>”:
    - i) Without Feature S233, “Multiset locators”, conforming SQL language shall not contain a <PL/I multiset locator variable>.
- 148) Specifications for Feature S241, “Transform functions”:
- a) Subclause 6.4, “<value specification> and <target specification>”:
    - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain CURRENT\_DEFAULT\_TRANSFORM\_GROUP.



- ii) Without Feature S241, “Transform functions”, conforming SQL language shall not contain CURRENT\_TRANSFORM\_GROUP\_FOR\_TYPE.
- b) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <transform group specification>.
- c) Subclause 11.58, “<transform definition>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <transform definition>.
- d) Subclause 11.62, “<drop transform statement>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <drop transform statement>.
- e) Subclause 13.1, “<SQL-client module definition>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <module transform group specification>.
- f) Subclause 19.9, “<set transform group statement>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <set transform group statement>.
- g) Subclause 21.1, “<embedded SQL host program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <embedded transform group specification>.
- h) Subclause 21.3, “<embedded SQL Ada program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain an <Ada user-defined type variable>.
- i) Subclause 21.4, “<embedded SQL C program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <C user-defined type variable>.
- j) Subclause 21.5, “<embedded SQL COBOL program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <COBOL user-defined type variable>.
- k) Subclause 21.6, “<embedded SQL Fortran program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <Fortran user-defined type variable>.
- l) Subclause 21.7, “<embedded SQL MUMPS program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <MUMPS user-defined type variable>.
- m) Subclause 21.8, “<embedded SQL Pascal program>”:

- i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <Pascal user-defined type variable>.
- n) Subclause 21.9, “<embedded SQL PL/I program>”:
  - i) Without Feature S241, “Transform functions”, conforming SQL language shall not contain a <PL/I user-defined type variable>.
- 149) Specifications for Feature S242, “Alter transform statement”:
  - a) Subclause 11.59, “<alter transform statement>”:
    - i) Without Feature S242, “Alter transform statement”, conforming SQL language shall not contain an <alter transform statement>.
- 150) Specifications for Feature S251, “User-defined orderings”:
  - a) Subclause 11.56, “<user-defined ordering definition>”:
    - i) Without Feature S251, “User-defined orderings”, conforming SQL shall not contain a <user-defined ordering definition>.

NOTE 491 — If MAP is specified, then the Conformance Rules of Subclause 9.9, “Equality operations”, apply. If ORDER FULL BY MAP is specified, then the Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.
  - b) Subclause 11.57, “<drop user-defined ordering statement>”:
    - i) Without Feature S251, “User-defined orderings”, conforming SQL language shall not contain a <drop user-defined ordering statement>.
- 151) Specifications for Feature S261, “Specific type method”:
  - a) Subclause 6.29, “<string value function>”:
    - i) Without Feature S261, “Specific type method”, conforming SQL language shall not contain a <specific type method>.
- 152) Specifications for Feature S271, “Basic multiset support”:
  - a) Subclause 6.1, “<data type>”:
    - i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset type>.
  - b) Subclause 6.5, “<contextually typed value specification>”:
    - i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain an <empty specification> that simply contains MULTISSET.
  - c) Subclause 6.24, “<multiset element reference>”:
    - i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset element reference>.
  - d) Subclause 6.27, “<numeric value function>”:
    - i) Without Feature S091, “Basic array support”, or Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <cardinality expression>.

## e) Subclause 6.38, “&lt;multiset value function&gt;”:

- i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset value function>.

NOTE 492 — The Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

## f) Subclause 6.39, “&lt;multiset value constructor&gt;”:

- i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <multiset value constructor>.

## g) Subclause 7.6, “&lt;table reference&gt;”:

- i) Without Feature S091, “Basic array support”, or Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <collection derived table>.

## h) Subclause 8.16, “&lt;member predicate&gt;”:

- i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <member predicate>.

NOTE 493 — The Conformance Rules of Subclause 9.9, “Equality operations”, also apply.

## i) Subclause 8.18, “&lt;set predicate&gt;”:

- i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <set predicate>.

NOTE 494 — The Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

## j) Subclause 10.9, “&lt;aggregate function&gt;”:

- i) Without Feature S271, “Basic multiset support”, conforming SQL language shall not contain a <computational operation> that immediately contains COLLECT.

## 153) Specifications for Feature S272, “Multisets of user-defined types”:

## a) Subclause 6.1, “&lt;data type&gt;”:

- i) Without Feature S272, “Multisets of user-defined types”, conforming SQL language shall not contain a <multiset type> that is based on a <data type> that contains a <path-resolved user-defined type name>.

## 154) Specifications for Feature S274, “Multisets of reference types”:

## a) Subclause 6.1, “&lt;data type&gt;”:

- i) Without Feature S274, “Multisets of reference types”, conforming SQL language shall not contain a <multiset type> that is based on a <data type> that contains a <reference type>.

## 155) Specifications for Feature S275, “Advanced multiset support”:

## a) Subclause 6.37, “&lt;multiset value expression&gt;”:

- i) Without Feature S275, “Advanced multiset support”, conforming SQL language shall not contain MULTiset UNION, MULTiset INTERSECTION, or MULTiset EXCEPT.

NOTE 495 — If MULTiset UNION DISTINCT, MULTiset INTERSECTION, or MULTiset EXCEPT is specified, then the Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

b) Subclause 8.17, “<submultiset predicate>”:

- i) Without Feature S275, “Advanced multiset support”, conforming SQL language shall not contain a <submultiset predicate>.

NOTE 496 — The Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

c) Subclause 9.9, “Equality operations”:

- i) Without Feature S275, “Advanced multiset support”, in conforming SQL language, the declared type of an operand of an equality operation shall not be multiset-ordered.

NOTE 497 — If the declared type of an operand *OP* of an equality operation is a multiset type, then *OP* is a multiset operand of a multiset element grouping operation. The Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, apply.

d) Subclause 10.9, “<aggregate function>”:

- i) Without Feature S275, “Advanced multiset support”, conforming SQL language shall not contain a <computational operation> that immediately contains FUSION or INTERSECTION.

NOTE 498 — If INTERSECTION is specified, then the Conformance Rules of Subclause 9.11, “Multiset element grouping operations”, also apply.

156) Specifications for Feature S281, “Nested collection types”:

a) Subclause 6.1, “<data type>”:

- i) Without Feature S281, “Nested collection types”, conforming SQL language shall not contain a collection type that is based on a <data type> that contains a <collection type>.

157) Specifications for Feature S291, “Unique constraint on entire row”:

a) Subclause 11.7, “<unique constraint definition>”:

- i) Without Feature S291, “Unique constraint on entire row”, conforming SQL language shall not contain UNIQUE(VALUE).

158) Specifications for Feature S301, “Enhanced UNNEST”:

a) Subclause 7.6, “<table reference>”:

- i) Without Feature S301, “Enhanced UNNEST”, in conforming SQL language, a <collection derived table> shall not simply contain more than one <collection value expression>.

159) Specifications for Feature T021, “BINARY and VARBINARY data types”:

a) Subclause 5.3, “<literal>”:

- i) Without Feature T041, “Basic LOB data type support”, or Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary string literal>.

b) Subclause 6.1, “<data type>”:

- i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary string type> that is not a <binary large object string type>.

c) Subclause 21.3, “<embedded SQL Ada program>”:

- i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain an <Ada BINARY variable>.

- ii) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain an <Ada VARBINARY variable>.
  - d) Subclause 21.4, “<embedded SQL C program>”:
    - i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <C BINARY variable>.
    - ii) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <C VARBINARY variable>.
  - e) Subclause 21.5, “<embedded SQL COBOL program>”:
    - i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <COBOL BINARY variable>.
  - f) Subclause 21.6, “<embedded SQL Fortran program>”:
    - i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <Fortran BINARY variable>.
    - ii) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <Fortran VARBINARY variable>.
  - g) Subclause 21.7, “<embedded SQL MUMPS program>”:
    - i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <MUMPS VARBINARY variable>.
  - h) Subclause 21.8, “<embedded SQL Pascal program>”:
    - i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <Pascal BINARY variable>.
  - i) Subclause 21.9, “<embedded SQL PL/I program>”:
    - i) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <PL/I BINARY variable>.
    - ii) Without Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <PL/I VARBINARY variable>.
- 160) Specifications for Feature T022, “Advanced support for BINARY and VARBINARY data types”:
- a) Subclause 6.29, “<string value function>”:
    - i) Without Feature T042, “Extended LOB data type support”, or Feature T022, “Advanced support for BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary value function>.
  - b) Subclause 8.5, “<like predicate>”:
    - i) Without Feature T042, “Extended LOB data type support”, or Feature T022, “Advanced support for BINARY and VARBINARY data types”, conforming SQL language shall not contain an <octet like predicate>.
- 161) Specifications for Feature T023, “Compound binary literals”:
- a) Subclause 5.3, “<literal>”:

- i) Without Feature T023, “Compound binary literals”, in conforming SQL language, a <binary string literal> shall contain exactly one repetition of “<quote> [ { <hexit> <hexit> }... ] <quote>”.

162) Specifications for Feature T024, “Spaces in binary literals”:

a) Subclause 5.3, “<literal>”:

- i) Without Feature T024, “Spaces in binary literals”, in conforming SQL language, a <binary string literal> shall not contain a <space>.

163) Specifications for Feature T031, “BOOLEAN data type”:

a) Subclause 5.3, “<literal>”:

- i) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <boolean literal>.

b) Subclause 6.1, “<data type>”:

- i) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <boolean type>.

c) Subclause 6.25, “<value expression>”:

- i) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <value expression> that is a <boolean value expression>.

d) Subclause 6.34, “<boolean value expression>”:

- i) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <boolean primary> that simply contains a <nonparenthesized value expression primary>.

e) Subclause 7.1, “<row value constructor>”:

- i) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <row value constructor predicand> that immediately contains a <boolean predicand>.

f) Subclause 10.9, “<aggregate function>”:

- i) Without Feature T031, “BOOLEAN data type”, conforming SQL language shall not contain a <computational operation> that immediately contains EVERY, ANY, or SOME.

164) Specifications for Feature T041, “Basic LOB data type support”:

a) Subclause 5.3, “<literal>”:

- i) Without Feature T041, “Basic LOB data type support”, or Feature T021, “BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary string literal>.

b) Subclause 6.1, “<data type>”:

- i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <binary large object string type>, a <character large object type>, or a <national character large object type>.

c) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <parameter type> that contains a <locator indication> and that simply contains a <data type> that identifies a large object type.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <returns data type> that contains a <locator indication> and that simply contains a <data type> that identifies a large object type.
- d) Subclause 21.3, “<embedded SQL Ada program>”:
- i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada CLOB variable>.
  - iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada BLOB locator variable>.
  - iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain an <Ada CLOB locator variable>.
- e) Subclause 21.4, “<embedded SQL C program>”:
- i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C CLOB variable>.
  - iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C BLOB locator variable>.
  - iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <C CLOB locator variable>.
- f) Subclause 21.5, “<embedded SQL COBOL program>”:
- i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL CLOB variable>.
  - iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL BLOB locator variable>.
  - iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <COBOL CLOB locator variable>.
- g) Subclause 21.6, “<embedded SQL Fortran program>”:
- i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran CLOB variable>.



- iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran BLOB locator variable>.
- iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Fortran CLOB locator variable>.
- h) Subclause 21.7, “<embedded SQL MUMPS program>”:
  - i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS CLOB variable>.
  - iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS BLOB locator variable>.
  - iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <MUMPS CLOB locator variable>.
- i) Subclause 21.8, “<embedded SQL Pascal program>”:
  - i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal CLOB variable>.
  - iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal BLOB locator variable>.
  - iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <Pascal BLOB variable>, <Pascal CLOB variable>, <Pascal CLOB locator variable>.
- j) Subclause 21.9, “<embedded SQL PL/I program>”:
  - i) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I BLOB variable>.
  - ii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I CLOB variable>.
  - iii) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I BLOB locator variable>.
  - iv) Without Feature T041, “Basic LOB data type support”, conforming SQL language shall not contain a <PL/I CLOB locator variable>.

165) Specifications for Feature T042, “Extended LOB data type support”:

- a) Subclause 6.12, “<cast specification>”:
  - i) Without Feature T042, “Extended LOB data type support”, conforming SQL language shall not contain a <cast operand> whose declared type is BINARY LARGE OBJECT or CHARACTER LARGE OBJECT.
  - ii) Without Feature T042, “Extended LOB data type support”, conforming SQL language shall not contain a <cast operand> whose declared type is NATIONAL CHARACTER LARGE OBJECT.



## b) Subclause 6.29, “&lt;string value function&gt;”:

- i) Without Feature T042, “Extended LOB data type support”, or Feature T022, “Advanced support for BINARY and VARBINARY data types”, conforming SQL language shall not contain a <binary value function>.

## c) Subclause 8.5, “&lt;like predicate&gt;”:

- i) Without Feature T042, “Extended LOB data type support”, or Feature T022, “Advanced support for BINARY and VARBINARY data types”, conforming SQL language shall not contain an <octet like predicate>.
- ii) Without Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <like predicate> shall not be of declared type CHARACTER LARGE OBJECT
- iii) Without Feature F421, “National character”, and Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <like predicate> shall not be of declared type NATIONAL CHARACTER LARGE OBJECT.

## d) Subclause 8.6, “&lt;similar predicate&gt;”:

- i) Without Feature T042, “Extended LOB data type support”, in conforming SQL language, a <character value expression> simply contained in a <similar predicate> shall not be of declared type CHARACTER LARGE OBJECT.

## e) Subclause 9.9, “Equality operations”:

- i) Without Feature T042, “Extended LOB data type support”, in conforming SQL language, the declared type of an operand of an equality operation shall not be LOB-ordered.

## 166) Specifications for Feature T043, “Multiplier T”:

## a) Subclause 5.2, “&lt;token&gt; and &lt;separator&gt;”:

- i) Without Feature T043, “Multiplier T”, in conforming SQL language, a <multiplier> shall not be T.

## 167) Specifications for Feature T044, “Multiplier P”:

## a) Subclause 5.2, “&lt;token&gt; and &lt;separator&gt;”:

- i) Without Feature T044, “Multiplier P”, in conforming SQL language, a <multiplier> shall not be P.

## 168) Specifications for Feature T051, “Row types”:

## a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain a <field name>.

## b) Subclause 6.1, “&lt;data type&gt;”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain a <row type>.

## c) Subclause 6.2, “&lt;field definition&gt;”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain a <field definition>.

d) Subclause 6.14, “<field reference>”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain a <field reference>.

e) Subclause 7.1, “<row value constructor>”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain an <explicit row value constructor> that immediately contains ROW.
- ii) Without Feature T051, “Row types”, conforming SQL language shall not contain a <contextually typed row value constructor> that immediately contains ROW.

f) Subclause 7.2, “<row value expression>”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain a <row value special case>.

g) Subclause 7.12, “<query specification>”:

- i) Without Feature T051, “Row types”, conforming SQL language shall not contain an <all fields reference>.

169) Specifications for Feature T052, “MAX and MIN for row types”:

a) Subclause 10.9, “<aggregate function>”:

- i) Without Feature T052, “MAX and MIN for row types”, conforming SQL language shall not contain a <computational operation> that immediately contains MAX or MIN in which the declared type of the <value expression> is a row type.

NOTE 499 — If DISTINCT is specified, then the Conformance Rules of Subclause 9.10, “Grouping operations”, also apply. If MAX or MIN is specified, then the Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.

170) Specifications for Feature T053, “Explicit aliases for all-fields reference”:

a) Subclause 7.12, “<query specification>”:

- i) Without Feature T053, “Explicit aliases for all-fields reference”, conforming SQL language shall not contain an <all fields column name list>.

NOTE 500 — If a <set quantifier> DISTINCT is specified, then the Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

171) Specifications for Feature T061, “UCS support”:

a) Subclause 6.1, “<data type>”:

- i) Without Feature T061, “UCS support”, conforming SQL language shall not contain a <char length units>.

b) Subclause 6.29, “<string value function>”:

- i) Without Feature T061, “UCS support”, conforming SQL language shall not contain a <normalize function>.

c) Subclause 8.12, “<normalized predicate>”:

- i) Without Feature T061, “UCS support”, conforming SQL language shall not contain a <normalized predicate>.

172) Specifications for Feature T071, “BIGINT data type”:

- a) Subclause 6.1, “<data type>”:
  - i) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain BIGINT.
- b) Subclause 21.3, “<embedded SQL Ada program>”:
  - i) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain an <Ada qualified type specification> that contains Interfaces.SQL.BIGINT.
  - ii) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain an <Ada unqualified type specification> that contains BIGINT.
- c) Subclause 21.4, “<embedded SQL C program>”:
  - i) Without Feature T071, “BIGINT data type”, conforming SQL language shall not contain a <C numeric variable> that contains long long.

173) Specifications for Feature T111, “Updatable joins, unions and columns”:

- a) Subclause 11.23, “<view definition>”:
  - i) Without Feature T111, “Updatable joins, unions and columns”, in conforming SQL language, if WITH CHECK OPTION is specified, then the viewed table shall be simply updatable.
- b) Subclause 14.3, “<cursor specification>”:
  - i) Without Feature T111, “Updatable joins, unions, and columns”, in conforming SQL language, if FOR UPDATE is specified, then *QE* shall be simply updatable.
- c) Subclause 14.9, “<delete statement: searched>”:
  - i) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain a <delete statement: searched> that contains a <target table> that identifies a table that is not simply updatable.
- d) Subclause 14.11, “<insert statement>”:
  - i) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain an <insert statement> that contains an <insertion target> that identifies a table that is not simply updatable.
- e) Subclause 14.12, “<merge statement>”:
  - i) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain a <merge statement> that contains an <target table> that identifies a table that is not simply updatable.
- f) Subclause 14.14, “<update statement: searched>”:
  - i) Without Feature T111, “Updatable joins, unions, and columns”, conforming SQL language shall not contain an <update statement: searched> that contains a <target table> that identifies a table that is not simply updatable.

174) Specifications for Feature T121, “WITH (excluding RECURSIVE) in query expression”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature T121, “WITH (excluding RECURSIVE) in query expression”, conforming SQL language shall not contain a <query name>.

b) Subclause 7.6, “<table reference>”:

- i) Without Feature T121, “WITH (excluding RECURSIVE) in query expression”, conforming SQL language shall not contain a <query name>.

c) Subclause 7.13, “<query expression>”:

- i) Without Feature T121, “WITH (excluding RECURSIVE) in query expression”, in conforming SQL language, a <query expression> shall not contain a <with clause>.

175) Specifications for Feature T122, “WITH (excluding RECURSIVE) in subquery”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature T122, “WITH (excluding RECURSIVE) in subquery”, in conforming SQL language, a <query expression> contained in a <query expression> shall not contain a <with clause>.

176) Specifications for Feature T131, “Recursive query”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature T131, “Recursive query”, conforming SQL language shall not contain a <query expression> that contains RECURSIVE.

b) Subclause 11.23, “<view definition>”:

- i) Without Feature T131, “Recursive query”, conforming SQL language shall not contain a <view definition> that immediately contains RECURSIVE.

177) Specifications for Feature T132, “Recursive query in subquery”:

a) Subclause 7.13, “<query expression>”:

- i) Without Feature T132, “Recursive query in subquery”, in conforming SQL language, a <query expression> contained in a <query expression> shall not contain RECURSIVE.

178) Specifications for Feature T141, “SIMILAR predicate”:

a) Subclause 8.6, “<similar predicate>”:

- i) Without Feature T141, “SIMILAR predicate”, conforming SQL language shall not contain a <similar predicate>.

179) Specifications for Feature T151, “DISTINCT predicate”:

a) Subclause 8.15, “<distinct predicate>”:

- i) Without Feature T151, “DISTINCT predicate”, conforming SQL language shall not contain a <distinct predicate>.

NOTE 501 — The Conformance Rules of Subclause 9.9, “Equality operations”, also apply.

180) Specifications for Feature T152, “DISTINCT predicate with negation”:

a) Subclause 8.15, “<distinct predicate>”:

- i) Without Feature T152, “DISTINCT predicate with negation”, conforming SQL language shall not contain a <distinct predicate part 2> that immediately contains NOT.

181) Specifications for Feature T171, “LIKE clause in table definition”:

a) Subclause 11.3, “<table definition>”:

- i) Without Feature T171, “LIKE clause in table definition”, conforming SQL language shall not contain a <like clause>.

182) Specifications for Feature T172, “AS subquery clause in table definition”:

a) Subclause 11.3, “<table definition>”:

- i) Without Feature T172, “AS subquery clause in table definition”, conforming SQL language shall not contain an <as subquery clause>.

183) Specifications for Feature T173, “Extended LIKE clause in table definition”:

a) Subclause 11.3, “<table definition>”:

- i) Without Feature T173, “Extended LIKE clause in table definition”, a <like clause> shall not contain <like options>.

184) Specifications for Feature T174, “Identity columns”:

a) Subclause 11.4, “<column definition>”:

- i) Without Feature T174, “Identity columns”, conforming SQL language shall not contain an <identity column specification>.

b) Subclause 11.18, “<alter identity column specification>”:

- i) Without Feature T174, “Identity columns”, an <alter column definition> shall not contain an <alter identity column specification>.

185) Specifications for Feature T175, “Generated columns”:

a) Subclause 11.4, “<column definition>”:

- i) Without Feature T175, “Generated columns”, conforming SQL language shall not contain a <generation clause>.

186) Specifications for Feature T176, “Sequence generator support”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <sequence generator name>.

b) Subclause 6.13, “<next value expression>”:

- i) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <next value expression>.

c) Subclause 11.63, “<sequence generator definition>”:

- i) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <sequence generator definition>.
- d) Subclause 11.64, “<alter sequence generator statement>”:
  - i) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain an <alter sequence generator statement>.
- e) Subclause 11.65, “<drop sequence generator statement>”:
  - i) Without Feature T176, “Sequence generator support”, conforming SQL language shall not contain a <drop sequence generator statement>.

187) Specifications for Feature T191, “Referential action RESTRICT”:

- a) Subclause 11.8, “<referential constraint definition>”:
  - i) Without Feature T191, “Referential action RESTRICT”, conforming SQL language shall not contain a <referential action> that contains RESTRICT.

188) Specifications for Feature T201, “Comparable data types for referential constraints”:

- a) Subclause 11.8, “<referential constraint definition>”:
  - i) Without Feature T201, “Comparable data types for referential constraints”, conforming SQL language shall not contain a <referencing columns> in which the data type of each referencing column is not the same as the data type of the corresponding referenced column.

NOTE 502 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

189) Specifications for Feature T211, “Basic trigger capability”:

- a) Subclause 7.6, “<table reference>”:
  - i) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain a <transition table name>.
- b) Subclause 11.40, “<trigger definition>”:
  - i) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain a <trigger definition>.
- c) Subclause 11.41, “<drop trigger statement>”:
  - i) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain a <drop trigger statement>.
- d) Subclause 12.3, “<privileges>”:
  - i) Without Feature T211, “Basic trigger capability”, conforming SQL language shall not contain an <action> that contains TRIGGER.

190) Specifications for Feature T212, “Enhanced trigger capability”:

- a) Subclause 11.40, “<trigger definition>”:
  - i) Without Feature T212, “Enhanced trigger capability”, in conforming SQL language, a <triggered action> shall contain FOR EACH ROW.

191) Specifications for Feature T213, “INSTEAD OF triggers”:

a) Subclause 11.40, “<trigger definition>”:

- i) Without Feature T213, “INSTEAD OF triggers”, in conforming SQL language, a <trigger action time> shall not immediately contain INSTEAD OF.

192) Specifications for Feature T231, “Sensitive cursors”:

a) Subclause 14.2, “<cursor properties>”:

- i) Without Feature T231, “Sensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains SENSITIVE.
- ii) Without Feature F791, “Insensitive cursors”, or Feature T231, “Sensitive cursors”, conforming SQL language shall not contain a <cursor sensitivity> that immediately contains ASENSITIVE.

193) Specifications for Feature T241, “START TRANSACTION statement”:

a) Subclause 17.1, “<start transaction statement>”:

- i) Without Feature T241, “START TRANSACTION statement”, conforming SQL language shall not contain a <start transaction statement>.

194) Specifications for Feature T251, “SET TRANSACTION statement: LOCAL option”:

a) Subclause 17.2, “<set transaction statement>”:

- i) Without Feature T251, “SET TRANSACTION statement: LOCAL option”, conforming SQL language shall not contain a <set transaction statement> that immediately contains LOCAL.

195) Specifications for Feature T261, “Chained transactions”:

a) Subclause 17.7, “<commit statement>”:

- i) Without Feature T261, “Chained transactions”, conforming SQL language shall not contain a <commit statement> that immediately contains CHAIN.

b) Subclause 17.8, “<rollback statement>”:

- i) Without Feature T261, “Chained transactions”, conforming SQL language shall not contain a <rollback statement> that immediately contains CHAIN.

196) Specifications for Feature T271, “Savepoints”:

a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <savepoint name>.

b) Subclause 17.5, “<savepoint statement>”:

- i) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <savepoint statement>.

c) Subclause 17.6, “<release savepoint statement>”:

- i) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <release savepoint statement>.

d) Subclause 17.8, “<rollback statement>”:



- i) Without Feature T271, “Savepoints”, conforming SQL language shall not contain a <savepoint clause>.

197) Specifications for Feature T272, “Enhanced savepoint management”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T272, “Enhanced savepoint management”, conforming SQL language shall not contain a <routine characteristics> that contains a <savepoint level indication>.

198) Specifications for Feature T281, “SELECT privilege with column granularity”:

a) Subclause 12.3, “<privileges>”:

- i) Without Feature T281, “SELECT privilege with column granularity”, in conforming SQL language, an <action> that contains SELECT shall not contain a <privilege column list>.

199) Specifications for Feature T285, “Enhanced derived column names”:

a) Subclause 7.12, “<query specification>”:

- i) Without Feature T285, “Enhanced derived column names”, in conforming SQL language, if any <derived column> in a <select list> does not specify an <as clause> and the <value expression> of that <derived column> is not a single column reference, then the <column name> of that column is implementation-dependent.

200) Specifications for Feature T301, “Functional dependencies”:

a) Subclause 7.10, “<having clause>”:

- i) Without Feature T301, “Functional dependencies”, in conforming SQL language, each column reference directly contained in the <search condition> shall be one of the following:
  - 1) An unambiguous reference to a grouping column of *T*.
  - 2) An outer reference.
- ii) Without Feature T301, “Functional dependencies”, in conforming SQL language, each column reference contained in a <query expression> in the <search condition> that references a column of *T* shall be one of the following:
  - 1) An unambiguous reference to a grouping column of *T*.
  - 2) Contained in an aggregated argument of a <set function specification>.

b) Subclause 7.11, “<window clause>”:

- i) Without Feature T301, “Functional dependencies”, in conforming SQL language, if *T* is a grouped table, then each column reference contained in <window clause> that references a column of *T* shall be a reference to a grouping column of *T* or be contained in an aggregated argument of a <set function specification>.

c) Subclause 7.12, “<query specification>”:

- i) Without Feature T301, “Functional dependencies”, in conforming SQL language, if *T* is a grouped table, then in each <value expression> contained in the <select list>, each <column reference> that references a column of *T* shall reference a grouping column or be specified in an aggregated argument of a <set function specification>.



d) Subclause 20.4, “<get descriptor statement>”:

- i) Without Feature T301, “Functional dependencies”, conforming SQL language shall not contain a <descriptor item name> that contains KEY\_MEMBER.

20l) Specifications for Feature T312, “OVERLAY function”:

a) Subclause 6.29, “<string value function>”:

- i) Without Feature T312, “OVERLAY function”, conforming SQL language shall not contain a <character overlay function>.
- ii) Without Feature T312, “OVERLAY function”, conforming SQL language shall not contain a <binary overlay function>.

202) Specifications for Feature T322, “Overloading of SQL-invoked functions and procedures”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T322, “Overloading of SQL-invoked functions and procedures”, conforming SQL language shall not contain a <schema routine> in which the schema identified by the explicit or implicit schema name of the <schema qualified routine name> includes a routine descriptor whose routine name is <schema qualified routine name>.

203) Specifications for Feature T323, “Explicit security for external routines”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T323, “Explicit security for external routines”, conforming SQL language shall not contain an <external security clause>.

204) Specifications for Feature T324, “Explicit security for SQL routines”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T324, “Explicit security for SQL routines”, conforming SQL language shall not contain a <rights clause>.

205) Specifications for Feature T325, “Qualified SQL parameter references”:

a) Subclause 6.6, “<identifier chain>”:

- i) Without Feature T325, “Qualified SQL parameter references”, conforming SQL language shall not contain an SQL parameter reference whose first <identifier> is the <qualified identifier> of a <routine name>.

b) Subclause 7.12, “<query specification>”:

- i) Without Feature T325, “Qualified SQL parameter references”, conforming SQL language shall not contain an <asterisked identifier chain> whose referent is an SQL parameter and whose first <identifier> is the <qualified identifier> of a <routine name>.

206) Specifications for Feature T326, “Table functions”:

a) Subclause 6.39, “<multiset value constructor>”:

- i) Without Feature T326, “Table functions”, a <multiset value constructor> shall not contain a <table value constructor by query>.

- b) Subclause 7.6, “<table reference>”:
  - i) Without Feature T326, “Table functions”, conforming SQL language shall not contain a <table function derived table>.
- c) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature T326, “Table functions”, conforming SQL language shall not contain a <returns table type>.

207) Specifications for Feature T331, “Basic roles”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <role name>.
- b) Subclause 12.4, “<role definition>”:
  - i) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <role definition>.
- c) Subclause 12.5, “<grant role statement>”:
  - i) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <grant role statement>.
- d) Subclause 12.6, “<drop role statement>”:
  - i) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <drop role statement>.
- e) Subclause 12.7, “<revoke statement>”:
  - i) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <revoke role statement>.
- f) Subclause 19.3, “<set role statement>”:
  - i) Without Feature T331, “Basic roles”, conforming SQL language shall not contain a <set role statement>.

208) Specifications for Feature T332, “Extended roles”:

- a) Subclause 6.4, “<value specification> and <target specification>”:
  - i) Without Feature T332, “Extended roles”, conforming SQL language shall not contain CURRENT\_ROLE.
- b) Subclause 11.1, “<schema definition>”:
  - i) Without Feature T332, “Extended roles”, in conforming SQL language a <schema authorization identifier> shall not be a <role name>.
- c) Subclause 11.5, “<default clause>”:
  - i) Without Feature T332, “Extended roles”, conforming SQL language shall not contain a <default option> that contains CURRENT\_ROLE.
- d) Subclause 12.3, “<privileges>”:

- i) Without Feature T332, “Extended roles”, conforming SQL language shall not contain a <grantor>.
- e) Subclause 12.4, “<role definition>”:
  - i) Without Feature T332, “Extended roles”, conforming SQL language shall not contain a <role definition> that immediately contains WITH ADMIN.
- f) Subclause 12.8, “Grantor determination”:
  - i) Without Feature T332, “Extended roles”, conforming SQL language shall contain no <grantor>.

209) Specifications for Feature T351, “Bracketed comments”:

- a) Subclause 5.2, “<token> and <separator>”:
  - i) Without Feature T351, “Bracketed comments”, conforming SQL language shall not contain a <bracketed comment>.

210) Specifications for Feature T431, “Extended grouping capabilities”:

- a) Subclause 6.9, “<set function specification>”:
  - i) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <grouping operation>.
- b) Subclause 7.9, “<group by clause>”:
  - i) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <rollup list>.
  - ii) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <cube list>.
  - iii) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain a <grouping sets specification>.
  - iv) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain an <empty grouping set>.
  - v) Without Feature T431, “Extended grouping capabilities”, conforming SQL language shall not contain an <ordinary grouping set> that contains a <grouping column reference list>.

211) Specifications for Feature T432, “Nested and concatenated GROUPING SETS”:

- a) Subclause 7.9, “<group by clause>”:
  - i) Without Feature T432, “Nested and concatenated GROUPING SETS”, conforming SQL language shall not contain a <grouping set list> that contains a <grouping sets specification>.
  - ii) Without Feature T432, “Nested and concatenated GROUPING SETS”, conforming SQL language shall not contain a <group by clause> that simply contains a <grouping sets specification> GSS where GSS is not the only <grouping element> simply contained in the <group by clause>.

NOTE 503 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

212) Specifications for Feature T433, “Multiargument GROUPING function”:

- a) Subclause 6.9, “<set function specification>”:

- i) Without Feature T433, “Multiargument GROUPING function”, conforming SQL language shall not contain a <grouping operation> that contains more than one <column reference>.

213) Specifications for Feature T434, “GROUP BY DISTINCT”:

a) Subclause 7.9, “<group by clause>”:

- i) Without Feature T434, “GROUP BY DISTINCT”, conforming SQL language shall not contain a <group by clause> that simply contains a <set quantifier>.

214) Specifications for Feature T441, “ABS and MOD functions”:

a) Subclause 6.27, “<numeric value function>”:

- i) Without Feature T441, “ABS and MOD functions”, conforming language shall not contain an <absolute value expression>.
- ii) Without Feature T441, “ABS and MOD functions”, conforming language shall not contain a <modulus expression>.

215) Specifications for Feature T461, “Symmetric BETWEEN predicate”:

a) Subclause 8.3, “<between predicate>”:

- i) Without Feature T461, “Symmetric BETWEEN predicate”, conforming SQL language shall not contain SYMMETRIC or ASYMMETRIC.

NOTE 504 — Since <between predicate> is an ordering operation, the Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.

216) Specifications for Feature T471, “Result sets return value”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T471, “Result sets return value”, conforming SQL language shall not contain a <returned result sets characteristic>.

b) Subclause 14.2, “<cursor properties>”:

- i) Without Feature T471, “Result sets return value”, conforming SQL language shall not contain a <cursor returnability>.

c) Subclause 20.15, “<allocate cursor statement>”:

- i) Without Feature T471, “Result sets return value”, conforming SQL language shall not contain a <result set cursor>.

217) Specifications for Feature T491, “LATERAL derived table”:

a) Subclause 7.6, “<table reference>”:

- i) Without Feature T491, “LATERAL derived table”, conforming SQL language shall not contain a <lateral derived table>.

218) Specifications for Feature T501, “Enhanced EXISTS predicate”:

a) Subclause 8.10, “<exists predicate>”:

- i) Without Feature T501, “Enhanced EXISTS predicate”, conforming SQL language shall not contain an <exists predicate> that simply contains a <table subquery> in which the <select list>

of a <query specification> directly contained in the <table subquery> does not comprise either an <asterisk> or a single <derived column>.

219) Specifications for Feature T511, “Transaction counts”:

- a) Subclause 23.1, “<get diagnostics statement>”:
  - i) Without Feature T511, “Transaction counts”, conforming SQL language shall not contain a <statement information item name> that contains TRANSACTIONS\_COMMITTED, TRANSACTIONS\_ROLLED\_BACK, or TRANSACTION\_ACTIVE.

220) Specifications for Feature T551, “Optional key words for default syntax”:

- a) Subclause 7.13, “<query expression>”:
  - i) Without Feature T551, “Optional key words for default syntax”, conforming SQL language shall not contain UNION DISTINCT, EXCEPT DISTINCT, or INTERSECT DISTINCT.
- b) Subclause 14.2, “<cursor properties>”:
  - i) Without Feature T551, “Optional key words for default syntax”, conforming SQL language shall not contain a <cursor holdability> that immediately contains WITHOUT HOLD.

221) Specifications for Feature T561, “Holdable locators”:

- a) Subclause 14.17, “<free locator statement>”:
  - i) Without Feature T561, “Holdable locators”, conforming SQL language shall not contain a <free locator statement>.
- b) Subclause 14.18, “<hold locator statement>”:
  - i) Without Feature T561, “Holdable locators”, conforming SQL language shall not contain a <hold locator statement>.

222) Specifications for Feature T571, “Array-returning external SQL-invoked functions”:

- a) Subclause 11.42, “<user-defined type definition>”:
  - i) Without Feature T571, “Array-returning external SQL-invoked functions”, conforming SQL language shall not contain a <method specification> that contains a <returns clause> that satisfies either of the following conditions:
    - 1) A <result cast from type> is specified that simply contains an <array type> and does not contain a <locator indication>.
    - 2) A <result cast from type> is not specified and <returns data type> simply contains an <array type> and does not contain a <locator indication>.
- b) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature T571, “Array-returning external SQL-invoked functions”, conforming SQL language shall not contain an <SQL-invoked routine> that defines an array-returning external function.

223) Specifications for Feature T572, “Multiset-returning external SQL-invoked functions”:

- a) Subclause 11.42, “<user-defined type definition>”:

- i) Without Feature T572, “Multiset-returning external SQL-invoked functions”, conforming SQL language shall not contain a <method specification> that contains a <returns clause> that satisfies either of the following conditions:
  - 1) A <result cast from type> is specified that simply contains a <multiset type> and does not contain a <locator indication>.
  - 2) A <result cast from type> is not specified and <returns data type> simply contains a <multiset type> and does not contain a <locator indication>.
- b) Subclause 11.51, “<SQL-invoked routine>”:
  - i) Without Feature T572, “Multiset-returning external SQL-invoked functions”, conforming SQL language shall not contain an <SQL-invoked routine> that defines a multiset-returning external function.

224) Specifications for Feature T581, “Regular expression substring function”:

- a) Subclause 6.29, “<string value function>”:
  - i) Without Feature T581, “Regular expression substring function”, conforming SQL language shall not contain a <regular expression substring function>.

225) Specifications for Feature T591, “UNIQUE constraints of possibly null columns”:

- a) Subclause 11.7, “<unique constraint definition>”:
  - i) Without Feature T591, “UNIQUE constraints of possibly null columns”, in conforming SQL language, if UNIQUE is specified, then the <column definition> for each column whose <column name> is contained in the <unique column list> shall contain NOT NULL.

NOTE 505 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

226) Specifications for Feature T601, “Local cursor references”:

- a) Subclause 5.4, “Names and identifiers”:
  - i) Without Feature T601, “Local cursor references”, a <cursor name> shall not contain a <local qualifier>.

227) Specifications for Feature T611, “Elementary OLAP operations”:

- a) Subclause 6.10, “<window function>”:
  - i) Without Feature T611, “Elementary OLAP operations”, conforming SQL language shall not contain a <window function>.
- b) Subclause 7.11, “<window clause>”:
  - i) Without Feature T611, “Elementary OLAP operations”, conforming SQL language shall not contain a <window specification>.
- c) Subclause 10.10, “<sort specification list>”:
  - i) Without Feature T611, “Elementary OLAP operations”, conforming SQL language shall not contain a <null ordering>.

NOTE 506 — The Conformance Rules of Subclause 9.12, “Ordering operations”, also apply.

228) Specifications for Feature T612, “Advanced OLAP operations”:

## a) Subclause 5.4, “Names and identifiers”:

- i) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window name>.

## b) Subclause 6.10, “&lt;window function&gt;”:

- i) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window name>.
- ii) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain PERCENT\_RANK or CUME\_DIST.
- iii) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window function> that simply contains ROW\_NUMBER and immediately contains a <window name or specification> whose window structure descriptor does not contain a window ordering clause.

## c) Subclause 6.27, “&lt;numeric value function&gt;”:

- i) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <width bucket function>.

## d) Subclause 7.11, “&lt;window clause&gt;”:

- i) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window clause>.
- ii) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain an <existing window name>.
- iii) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <window frame exclusion>.

NOTE 507 — The Conformance Rules of Subclause 9.10, “Grouping operations”, also apply.

## e) Subclause 10.9, “&lt;aggregate function&gt;”:

- i) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <hypothetical set function> or an <inverse distribution function>.
- ii) Without Feature T612, “Advanced OLAP operations”, conforming SQL language shall not contain a <filter clause>.

## 229) Specifications for Feature T613, “Sampling”:

## a) Subclause 7.6, “&lt;table reference&gt;”:

- i) Without Feature T613, “Sampling”, conforming SQL language shall not contain a <sample clause>.

## 230) Specifications for Feature T621, “Enhanced numeric functions”:

## a) Subclause 6.27, “&lt;numeric value function&gt;”:

- i) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <natural logarithm>.



- ii) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain an <exponential function>.
  - iii) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <power function>.
  - iv) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <square root>.
  - v) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <floor function>.
  - vi) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <ceiling function>.
- b) Subclause 10.9, “<aggregate function>”:
- i) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <computational operation> that immediately contains STDDEV\_POP, STDDEV\_SAMP, VAR\_POP, or VAR\_SAMP.
  - ii) Without Feature T621, “Enhanced numeric functions”, conforming SQL language shall not contain a <binary set function type>.

23l) Specifications for Feature T641, “Multiple column assignment”:

- a) Subclause 14.15, “<set clause list>”:
- i) Without Feature T641, “Multiple column assignment”, conforming SQL language shall not contain a <multiple column assignment>.

232) Specifications for Feature T651, “SQL-schema statements in SQL routines”:

- a) Subclause 11.51, “<SQL-invoked routine>”:
- i) Without Feature T651, “SQL-schema statements in SQL routines”, conforming SQL language shall not contain an <SQL routine body> that contains an SQL-schema statement.

233) Specifications for Feature T652, “SQL-dynamic statements in SQL routines”:

- a) Subclause 11.51, “<SQL-invoked routine>”:
- i) Without Feature T652, “SQL-dynamic statements in SQL routines”, conforming SQL language shall not contain an <SQL routine body> that contains an SQL-dynamic statement.

234) Specifications for Feature T653, “SQL-schema statements in external routines”:

- a) Subclause 11.51, “<SQL-invoked routine>”:
- i) Without Feature T653, “SQL-schema statements in external routines”, conforming SQL language shall not contain an <external routine name> that identifies a program in which an SQL-schema statement appears.

235) Specifications for Feature T654, “SQL-dynamic statements in external routines”:

- a) Subclause 11.51, “<SQL-invoked routine>”:



- i) Without Feature T654, “SQL-dynamic statements in external routines”, conforming SQL language shall not contain an <external routine name> that identifies a program in which an SQL-dynamic statement appears.

236) Specifications for Feature T655, “Cyclically dependent routines”:

a) Subclause 11.51, “<SQL-invoked routine>”:

- i) Without Feature T655, “Cyclically dependent routines”, conforming SQL language shall not contain an <SQL routine body> that contains a <routine invocation> whose subject routine is generally dependent on the routine descriptor of the SQL-invoked routine specified by <SQL-invoked routine>.

*(Blank page)*

## Annex B

### (informative)

## Implementation-defined elements

*This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-3.  
 This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-4.  
 This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-9.  
 This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-10.  
 This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-11.  
 This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-13.  
 This Annex is modified by Annex B, “Implementation-defined elements”, in ISO/IEC 9075-14.*

This Annex references those features that are identified in the body of this part of ISO/IEC 9075 as implementation-defined.

- 1) Subclause 4.2.2, “Comparison of character strings”:
  - a) The specific character set associated with the subtype of character string represented by the <key word>s NATIONAL CHARACTER is implementation-defined.
  - b) The circumstances in which conversion of non-UCS character string expressions from one character set to another is automatic is implementation-defined.
- 2) Subclause 4.2.4, “Character repertoires”:
  - a) The character repertoires, including standard defined repertoires, supported by the SQL-implementation are implementation-defined.
  - b) The character set named SQL\_TEXT is an implementation-defined character set that contains every character that is in <SQL language character> and all characters that are in other character sets supported by the SQL-implementation.
  - c) The character set named SQL\_IDENTIFIER is an implementation-defined character set that contains every character that is in <SQL language character> and all characters that the SQL-implementation supports for use in <regular identifier>s, which is the same as the repertoire that the SQL-implementation supports for use in <delimited identifier>s.
- 3) Subclause 4.2.5, “Character encoding forms”:
  - a) The character encodings in the SQL\_CHARACTER character encoding form are implementation-defined.
  - b) The character encodings in the SQL\_TEXT character encoding form are implementation-defined.
  - c) The character encodings in the SQL\_IDENTIFIER character encoding form are implementation-defined.
  - d) Which character encoding forms, including standard defined encoding forms, is implemented are implementation-defined.

4) Subclause 4.2.6, “Collations”:

- a) The collations, including standard defined collations, supported by the SQL-implementation are implementation-defined.
- b) The ordering specified by the SQL\_CHARACTER collation is implementation-defined.
- c) The ordering specified by the SQL\_TEXT collation is implementation-defined.
- d) The ordering specified by the SQL\_IDENTIFIER collation is implementation-defined.

5) Subclause 4.2.7, “Character sets”:

- a) It is implementation-defined which collation, UCS\_BASIC or UNICODE, is the default for the UTF8, UTF16, and UTF32 character sets.

6) Subclause 4.2.8, “Universal character sets”:

- a) With the exception of <normalize function> and <normalized predicate>, the result of any operation on an unnormalized UCS string is implementation-defined.

7) Subclause 4.3.2, “Binary string comparison”:

- a) For binary string values other than binary large object string values, it is implementation-defined whether trailing X'00's are considered significant when comparing two binary string values that are otherwise equivalent.

8) Subclause 4.4, “Numbers”:

- a) Whether truncation or rounding is performed when trailing digits are removed from a numeric value is implementation-defined.
- b) When an approximation is obtained by truncation or rounding and there are more than one approximation, then it is implementation-defined which approximation is chosen.
- c) It is implementation-defined which numeric values have approximations obtained by rounding or truncation for a given approximate numeric type.
- d) The boundaries within which the normal rules of arithmetic apply are implementation-defined.

9) Subclause 4.4.2, “Characteristics of numbers”:

- a) When converting between numeric data types, if least significant digits are lost, then it is implementation-defined whether rounding or truncation occurs.

10) Subclause 4.6.2, “Datetimes”:

- a) Whether an SQL-implementation supports leap seconds, and the consequences of such support for date and interval arithmetic, are implementation-defined.

11) Subclause 4.9, “Reference types”:

- a) In a host variable, a reference type is materialized as an *n*-octet value, where *N* is implementation-defined.

12) Subclause 4.13, “Columns, fields, and attributes”:

- a) A column of a base table is known not null if there is some implementation-defined rule that enables the SQL-implementation to correctly infer that the value of the column can never be null.

## 13) Subclause 4.14.6, “Operations involving tables”:

- a) If a <table reference> contains a <sample clause>, and the <sample clause> contains <repeatable clause>, then repeated executions of that <table reference> return a result table with identical rows for a given <repeat argument>, provided certain implementation-defined conditions are satisfied.

## 14) Subclause 4.15.3, “Window functions”:

- a) If PERCENT\_RANK is specified, then the declared type of the result is approximate numeric with implementation-defined precision.
- b) If CUME\_DIST is specified, then the declared type of the result is approximate numeric with implementation-defined precision.

## 15) Subclause 4.17.3, “Table constraints”:

- a) The ordering of the lists of referencing column names and referenced column names in a referential constrict descriptor is implementation-defined, but shall be such that corresponding column names occupy corresponding positions in each list.

## 16) Subclause 4.18, “Functional dependencies”:

- a) An SQL-implementation may define additional known functional dependencies.

## 17) Subclause 4.22, “SQL-client modules”:

- a) The mechanisms by which SQL-client modules are created or destroyed are implementation-defined.
- b) The manner in which an association between an SQL-client module and an SQL-agent is defined is implementation-defined.
- c) Whether a compilation unit may invoke or transfer control to other compilation units, written in the same or a different programming language is implementation-defined.

## 18) Subclause 4.23, “Embedded syntax”:

- a) Whether portions of name spaces are reserved by an SQL-implementation for the names of procedures, subroutines, program variables, branch labels, <SQL-client module definition>s, <externally-invoked procedure>s, or other host language or host language environment elements is implementation-defined. If portions of the name spaces are so reserved, the portion or portions reserved are also implementation-defined.

## 19) Subclause 4.24, “Dynamic SQL concepts”:

- a) Within an SQL-session, all prepared statements belong to the same implementation-defined <SQL-client module definition> that is different from any other <SQL-client module definition> that exists simultaneously in the environment.

## 20) Subclause 4.25, “Direct invocation of SQL”:

- a) The method of invoking <direct SQL statement>s, the method of raising conditions as a result of <direct SQL statement>s, the method of accessing diagnostic information, and the method of returning the results are all implementation-defined.

## 21) Subclause 4.32, “Cursors”:

- a) If a sensitive or asensitive holdable cursor is held open for a subsequent SQL-transaction, then whether any significant changes made to SQL-data (by this or any subsequent SQL-transaction in which the

cursor is held open) will be visible through that cursor in the subsequent SQL-transaction before that cursor is closed is implementation-defined.

22) Subclause 4.32.1, “General description of cursors”:

- a) The operational returnability of a received cursor is implementation-defined.

23) Subclause 4.34, “Basic security model”:

- a) The mapping of <authorization identifier>s to operating system users is implementation-defined.

24) Subclause 4.35, “SQL-transactions”:

- a) It is implementation-defined whether or not the execution of an SQL-data statement is permitted to occur within the same SQL-transaction as the execution of an SQL-schema statement. If it does occur, then the effect on any open cursor or deferred constraint is also implementation-defined.
- b) If an SQL-implementation detects unrecoverable errors and implicitly initiates the execution of a <rollback statement>, an exception condition is raised with an *implementation-defined exception code*.
- c) It is implementation-defined whether or not the dynamic execution of an <SQL dynamic data statement> is permitted to occur within the same SQL-transaction as the dynamic execution of an SQL-schema statement. If it does occur, then the effect on any prepared dynamic statement is also implementation-defined.
- d) It is implementation-defined whether or not, or how, a <rollback statement> that references a <savepoint specifier> affects diagnostics area contents, the contents of SQL descriptor areas, and the status of prepared statements.

25) Subclause 4.36, “SQL-connections”:

- a) It is implementation-defined how an SQL-implementation uses <SQL-server name> to determine the location, identity, and communication protocol required to access the SQL-server and initiate an SQL-session.

26) Subclause 4.37, “SQL-sessions”:

- a) When an SQL-session is initiated other than through the use of an explicit <connect statement>, then an SQL-session associated with an implementation-defined SQL-server is initiated. The default SQL-server is implementation-defined.
- b) The mechanism and rules by which an SQL-implementation determines whether a call to an <externally-invoked procedure> is the last call within the last active SQL-client module is implementation-defined.
- c) An SQL-session uses one or more implementation-defined schemas that contain the instances of any global temporary tables, created local temporary tables, or declared local temporary tables within the SQL-session.
- d) When an SQL-session is initiated, there is an implementation-defined default time zone used as the current default time zone displacement of the SQL-session.
- e) When an SQL-session is initiated other than through the use of an explicit <connect statement>, there is an implementation-defined initial value of the SQL-session user identifier.
- f) When an SQL-session is initiated, there is an implementation-defined default catalog whose name is used to effectively qualify all unqualified <schema name>s contained in <preparable statement>s that

are dynamically prepared in the current SQL-session through the execution of `<prepare statement>s` and `<execute immediate statement>s`.

- g) When an SQL-session is initiated, there is an implementation-defined default schema whose name is used to effectively qualify all unqualified `<schema qualified name>s` contained in `<preparable statement>s` that are dynamically prepared in the current SQL-session through the execution of `<prepare statement>s` and `<execute immediate statement>s`.
- h) The value of the current SQL-path before a successful execution of `<set path statement>` is implementation-defined.

27) Subclause 5.1, “`<SQL terminal character>`”:

- a) The end-of-line indicator (`<newline>`) is implementation-defined.

28) Subclause 5.2, “`<token>` and `<separator>`”:

- a) Equivalence of two `<regular identifier>s`, or of a `<regular identifier>` and a `<delimited identifier>` is determined using an implementation-defined collation that is sensitive to case.
- b) When the source character set does not contain `<reverse solidus>`, the character used as the default `<Unicode escape character>` is implementation-defined.

29) Subclause 5.3, “`<literal>`”:

- a) The `<character set name>` character set of the used to represent national characters is implementation-defined.
- b) The declared type of an `<exact numeric literal>` is implementation-defined.
- c) The declared type of an `<approximate numeric literal>` is implementation-defined.
- d) The declared type of a `<binary string literal>` is implementation-defined.

30) Subclause 5.4, “Names and identifiers”:

- a) If a `<schema name>` contained in a `<schema name clause>` but not contained in an SQL-client module does not contain a `<catalog name>`, then an implementation-defined `<catalog name>` is implicit.
- b) If a `<schema name>` contained in a `<module authorization clause>` does not contain a `<catalog name>`, then an implementation-defined `<catalog name>` is implicit.
- c) Those `<identifier>s` that are valid `<authorization identifier>s` are implementation-defined.
- d) Those `<identifier>s` that are valid `<catalog name>s` are implementation-defined.
- e) All `<transcoding name>s` are implementation-defined.
- f) If a `<schema name>` contained in a `<preparable statement>` that is dynamically prepared in the current SQL-session through the execution of a `<prepare statement>` or an `<execute immediate statement>` does not contain a `<catalog name>`, then the implementation-defined `<catalog name>` for the SQL-session is implicit.
- g) If a `<schema qualified name>` contained in a `<preparable statement>` that is dynamically prepared in the current SQL-session through the execution of a `<prepare statement>` or an `<execute immediate statement>` does not contain a `<schema name>`, then the implementation-defined `<schema name>` for the SQL-session is implicit.

31) Subclause 6.1, “<data type>”:

- a) The <character set name> associated with NATIONAL CHARACTER is implementation-defined.
- b) If a <precision> is omitted, then an implementation-defined <precision> is implicit.
- c) The decimal precision of a data type defined as DECIMAL for each value specified by <precision> is implementation-defined.
- d) The precisions of data types defined as SMALLINT, INTEGER, and BIGINT are implementation-defined, but all three data types have the same radix.
- e) The binary precision of a data type defined as FLOAT for each value specified by <precision> is implementation-defined.
- f) The precision of a data type defined as REAL is implementation-defined.
- g) The precision of a data type defined as DOUBLE PRECISION is implementation-defined, but greater than that for REAL.
- h) For every <data type>, the limits of the <data type> are implementation-defined.
- i) The maximum lengths for character string types and binary string types are implementation-defined.
- j) If CHARACTER SET is not specified for <character string type>, then the character set is implementation-defined.
- k) For the <exact numeric type>s DECIMAL and NUMERIC, the maximum values of <precision> and of <scale> are implementation-defined.
- l) The transformation *ENNF()* of an <exact numeric type> to its normal form, to obtain the data type name saved in a numeric data type descriptor, is implementation-defined, though it shall adhere to the following constraints:
  - i) For every <exact numeric type> *ENT*, *ENNF(ENT)* shall not specify DEC or INT.  
 NOTE 508 — The preceding requirement prohibits the function *ENNF* from returning a value that uses the abbreviated spelling of the two data types; the function shall instead return the long versions of DECIMAL or INTEGER.
  - ii) For every <exact numeric type> *ENT*, the precision, scale, and radix of *ENNF(ENT)* shall be the precision, scale, and radix of *ENT*.
  - iii) For every <exact numeric type> *ENT*, *ENNF(ENT)* shall be the same as *ENNF(ENNF(ENT))*.
  - iv) For every <exact numeric type> *ENT*, if *ENNF(ENT)* specifies DECIMAL, then *ENNF(ENT)* shall specify <precision>, and the precision of *ENNF(ENT)* shall be the value of the <precision> specified in *ENNF(ENT)*.
- m) For the <approximate numeric type> FLOAT, the maximum value of <precision> is implementation-defined.
- n) The transformation *ANNF()* of an <approximate numeric type> to its normal form, to obtain the data type name saved in a numeric data type descriptor, is implementation-defined, though it shall adhere to the following constraints:
  - i) For every <approximate numeric type> *ANT*, the precision of *ANNF(ANT)* shall be the precision of *ANT*.



- ii) For every <approximate numeric type> *ANT*, *ANNF(ANT)* shall be the same as *ANNF(ANNF(ANT))*.
- iii) For every <approximate numeric type> *ANT*, if *ANNF(ANT)* specifies *FLOAT*, then *ANNF(ANT)* shall specify <precision>, and the precision of *ANNF(ANT)* shall be the value of the <precision> specified in *ANNF(ANT)*.
- o) For the <approximate numeric type>s *FLOAT*, *REAL*, and *DOUBLE PRECISION*, the maximum and minimum values of the exponent are implementation-defined.
- p) The maximum value of <time fractional seconds precision> is implementation-defined, but shall not be less than 6.
- q) The maximum values of <time precision> and <timestamp precision> for a <datetime type> are the same implementation-defined value.
- r) If the maximum cardinality of an <array type> is omitted, then an implementation-defined maximum cardinality is implicit.

32) Subclause 6.4, “<value specification> and <target specification>”:

- a) Whether the character string of the <value specification>s *CURRENT\_USER*, *SESSION\_USER*, and *SYSTEM\_USER* is variable-length or fixed-length, and its maximum length if it is variable-length or its length if it is fixed-length, are implementation-defined.
- b) Whether the character string of the <value specification>s *CURRENT\_CATALOG* and *CURRENT\_SCHEMA* is variable-length or fixed-length, and its maximum length if it is variable-length or its length if it is fixed-length, are implementation-defined.
- c) The value specified by *SYSTEM\_USER* is an implementation-defined string that represents the operating system user who executed the SQL-client module that contains the SQL-statement whose execution caused the *SYSTEM\_USER* <general value specification> to be evaluated.
- d) Whether the data type of *CURRENT\_PATH* is fixed-length or variable-length, and its length if it is fixed-length or its maximum length if it is variable-length, are implementation-defined.
- e) If a <target specification> or <simple target specification> is assigned a value that is a zero-length character string, then it is implementation-defined whether an exception condition is raised: *data exception — zero-length character string*.
- f) Without Feature F611, “Indicator data types”, in conforming SQL language, the declared types of <indicator parameter>s and <indicator variable>s shall be the same implementation-defined data type.

33) Subclause 6.9, “<set function specification>”:

- a) The precision of the value derived from application of the *COUNT* function is implementation-defined.
- b) The precision of the value derived from application of the *SUM* function to a declared type of exact numeric is implementation-defined.
- c) The precision and scale of the value derived from application of the *AVG* function to a declared type of exact numeric is implementation-defined.
- d) The precision of the value derived from application of the *SUM* function or *AVG* function to a data type of approximate numeric is implementation-defined.

34) Subclause 6.12, “<cast specification>”:

- a) Whether to round or truncate when casting to exact numeric, approximate numeric, datetime, or interval data types is implementation-defined.
- b) When casting from an exact numeric value to an interval type, if the scale of the exact numeric value is greater than the number of digits of fractional seconds precision of the interval type, it is implementation-defined whether the cast operation performs rounding or truncation to adjust the source value to the limitations of the target type.

35) Subclause 6.26, “<numeric value expression>”:

- a) When the declared type of both operands of the addition, subtraction, multiplication, or division operator is exact numeric, the declared type of the result is an implementation-defined exact numeric type.
- b) When the declared type of both operands of the division operator is exact numeric, the scale of the result is implementation-defined.
- c) When the declared type of either operand of an arithmetic operator is approximate numeric, the declared type of the result is an implementation-defined approximate numeric type.
- d) Whether to round or truncate when performing division is implementation-defined.

36) Subclause 6.27, “<numeric value function>”:

- a) The declared type of <position expression> is an implementation-defined exact numeric type with scale 0 (zero).
- b) The declared type of <extract expression> is an implementation-defined exact numeric type. If <primary datetime field> specifies SECOND, then the scale is implementation-defined; otherwise, the scale is 0 (zero).
- c) The declared type of <length expression> is an implementation-defined exact numeric type with scale 0 (zero).
- d) If <cardinality expression> is specified, then the declared type of the result is an implementation-defined exact numeric type with scale 0 (zero).
- e) The declared type of the result of <natural logarithm> is an implementation-defined approximate numeric.
- f) The declared type of the result of <exponential function> is an implementation-defined approximate numeric.
- g) The declared type of the result of <power function> is an implementation-defined approximate numeric.
- h) The declared types of the results of <floor function> and of <ceiling function> are implementation-defined exact numeric type with scale 0 (zero) if the declared type of the simply contained <numeric value expression> is exact numeric; otherwise, the declared types of the results are implementation-defined approximate numeric types.
- i) The declared type of <regex occurrences function> is an implementation-defined exact numeric type with scale 0 (zero).
- j) The declared type of <regex position expression> is an implementation-defined exact numeric type with scale 0 (zero).

37) Subclause 6.28, “<string value expression>”:

- a) If the result of the <character value expression> is a zero-length character string, then it is implementation-defined whether an exception condition is raised: *data exception — zero-length character string*.
- b) If the character encoding form of <character factor> is UTF8, UTF16, or UTF32, and either of the operands is not normalized, then the result is implementation-defined.

38) Subclause 6.29, “<string value function>”:

- a) The maximum length of <character transliteration> or <transcoding> is implementation-defined.
- b) The character set of the result of a <transcoding> is implementation-defined.
- c) If <normalize function result length> is not specified, then the maximum length of the declared type of the <normalize function> is implementation-defined.
- d) If no <XQuery replacement string> is specified, then the maximum length of the declared type of <regex transliteration> is implementation-defined.

39) Subclause 6.32, “<interval value expression>”:

- a) The difference of two values of type TIME (with or without time zone) is constrained to be between –24:00:00 and +24:00:00 (excluding each end point); it is implementation-defined which of two non-zero values in this range is the result, although the computation shall be deterministic.
- b) When an interval is produced from the difference of two datetimes, the choice of whether to round or truncate is implementation-defined.
- c) The result's <interval leading field precision> is implementation-defined, but shall not be less than the <interval leading field precision> of the <interval primary>.
- d) The <interval leading field precision> is implementation-defined, but shall be sufficient to represent all interval values with the interval fields and <interval leading field precision> of <interval value expression 1> as well as all interval values with the interval fields and <interval leading field precision> of <interval term 1>.

40) Subclause 6.34, “<boolean value expression>”:

- a) A <boolean value expression> is a known-not-null condition for *X* if *BVE* conforms to an implementation-defined rule that enables the SQL-implementation to correctly infer that, when *BVE* is *True*, then *X* cannot be null.

41) Subclause 7.7, “<joined table>”:

- a) A column *CR* of the result of the <joined table> is known not null if *CR* conforms to an implementation-defined rule that correctly infers that the value of *CR* cannot be null.

42) Subclause 7.12, “<query specification>”:

- a) An SQL-implementation may define additional implementation-defined rules for recognizing known-not-null columns.

43) Subclause 8.2, “<comparison predicate>”:

- a) For two binary string values *X* and *Y*, neither of which is a binary large object string value, it is implementation-defined whether *X* is less than *Y* or whether *X* is equal to *Y*, if the only difference between *X* and *Y* are trailing X'00's in *Y*.

- b) It is implementation-defined whether the syntactic transformations to define all comparison predicates in terms of = and < are applied in the case of user-defined types whose comparison category is RELATIVE.

44) Subclause 8.5, “<like predicate>”:

- a) It is implementation-defined which collations can be used as collations for the <like predicate>.

45) Subclause 8.6, “<similar predicate>”:

- a) It is implementation-defined which collations can be used as collations for the <similar predicate>.

46) Subclause 9.1, “Retrieval assignment”:

- a) If a value *V* is approximate numeric and a target *T* is exact numeric, then whether the approximation of *V* retrieved into *T* is obtained by rounding or by truncation is implementation-defined.
- b) If a value *V* is datetime with a greater precision than a target *T*, then it is implementation-defined whether the approximation of *V* retrieved into *T* is obtained by rounding or truncation.
- c) If a value *V* is interval with a greater precision than a target *T*, then it is implementation-defined whether the approximation of *V* retrieved into *T* is obtained by rounding or by truncation.

47) Subclause 9.2, “Store assignment”:

- a) If a value *V* is approximate numeric and a target *T* is exact numeric, then whether the approximation of *V* stored into *T* is obtained by rounding or by truncation is implementation-defined.
- b) If a value *V* is datetime with a greater precision than a target *T*, then it is implementation-defined whether the approximation of *V* stored into *T* is obtained by rounding or truncation.
- c) If a value *V* is interval with a greater precision than a target *T*, then it is implementation-defined whether the approximation of *V* stored into *T* is obtained by rounding or by truncation.

48) Subclause 9.3, “Result of data type combinations”:

- a) If all of the data types in *DTS* are exact numeric, then the result data type is exact numeric with implementation-defined precision.
- b) If any data type in *DTS* is approximate numeric, then each data type in *DTS* shall be numeric and the result data type is approximate numeric with implementation-defined precision.

49) Subclause 9.16, “XQuery regular expression matching”:

- a) If the character repertoire of the pattern or flag is not UCS, then the conversion to UCS is implementation-defined.
- b) If the character repertoire of the subject string is not UCS, then the definition of “XQuery regular expression match” is implementation-defined.

50) Subclause 9.17, “XQuery regular expression replacement”:

- a) If the character repertoire of the pattern, replacement string, or flag is not UCS, then the conversion to UCS is implementation-defined.
- b) If the character repertoire of the subject string is not UCS, then the result of the Subclause is implementation-defined.

51) Subclause 9.24, “Creation of a sequence generator”:

- a) If <sequence generator maxvalue option> specifies NO MAXVALUE or if <sequence generator max-value option> is not specified, then a <sequence generator max value> that is an implementation-defined <signed numeric literal> of declared type *DT* is implicit.
- b) If <sequence generator minvalue option> specifies NO MINVALUE or if <sequence generator minvalue option> is not specified, then a <sequence generator min value> that is an implementation-defined <signed numeric literal> of declared type *DT* is implicit.

52) Subclause 9.25, “Altering a sequence generator”:

- a) If <sequence generator maxvalue option> specifies NO MAXVALUE, then a <sequence generator max value> that is an implementation-defined <signed numeric literal> of declared type *DT* is implicit.
- b) If <sequence generator minvalue option> specifies NO MINVALUE, then a <sequence generator min value> that is an implementation-defined <signed numeric literal> of declared type *DT* is implicit.

53) Subclause 10.1, “<interval qualifier>”:

- a) The maximum value of <interval leading field precision> is implementation-defined, but shall not be less than 2.
- b) The maximum value of <interval fractional seconds precision> is implementation-defined, but shall not be less than 2.

54) Subclause 10.4, “<routine invocation>”:

- a) If  $p_i$  is an output SQL parameter, then  $CPV_i$  is an implementation-defined value of type  $T_i$ .
- b) Whether a syntax error occurs if an <SQL routine body> contains an <SQL connection statement> or an <SQL transaction statement> is implementation-defined.
- c) When a new SQL-session context *RSC* is created, the current default catalog name, current default unqualified schema name, the current default character set name, the SQL-path of the current SQL-session, the current default time zone displacement of the current SQL-session, and the contents of all SQL dynamic descriptor areas are set to implementation-defined values.
- d) If *R* is an external routine, then it is implementation-defined whether the identities of all instances of created local temporary tables that are referenced in the <SQL-client module definition> of *p*, declared local temporary tables that are defined by <temporary table declaration>s that are contained in the <SQL-client module definition> of *p*, and the cursor instance descriptors of all open cursors that are not global extended dynamic cursors and whose SQL-client module is the SQL-client module of *p* are removed from *RSC*.
- e) If *R* is an external routine, then the method and time of binding of *P* to the schema that includes *R* is implementation-defined.
- f) After the completion of *p*, it is implementation-defined whether open cursors declared in the <SQL-client module definition> of *p* are closed and destroyed, whether local temporary tables associated with *RCS* are destroyed, and whether prepared statements prepared by *p* are deallocated.
- g) If *R* is an SQL-invoked procedure, then for each SQL parameter that is an output SQL parameter or both an input and output SQL parameter whose corresponding argument was not assigned a value, that corresponding argument is set to an implementation-defined value of the appropriate type.

- h) If the external security characteristic of an external SQL-invoked routine is IMPLEMENTATION DEFINED, then the user identifier and role name in the first cell of the authorization stack of the new SQL-session context are implementation-defined.

55) Subclause 10.9, “<aggregate function>”:

- a) If COUNT is specified, then the declared type of the result is an implementation-defined exact numeric type with scale of 0 (zero).
- b) If SUM or AVG is specified, then:
  - i) If SUM is specified and the declared type of the argument is exact numeric with scale *S*, then the declared type of the result is an implementation-defined exact numeric type with scale *S*.
  - ii) If AVG is specified and the declared type *DT* of the argument is exact numeric, then the declared type of the result is an implementation-defined exact numeric type with precision not less than the precision of *DT* and scale not less than the scale of *DT*.
  - iii) If the declared type *DT* of the argument is approximate numeric, then the declared type of the result is an implementation-defined approximate numeric type with precision not less than the precision of *DT*.
- c) If VAR\_POP or VAR\_SAMP is specified, then the declared type of the result is an implementation-defined approximate numeric type. If the declared type of the argument is approximate numeric, then the precision of the result is not less than the precision of the argument.
- d) If <binary set function type> is specified, then

Case:

- i) If REGR\_COUNT is specified, then the declared type of the result is an implementation-defined exact numeric type with scale of 0 (zero).
  - ii) Otherwise, the declared type of the result is an implementation-defined approximate numeric type. If either argument is approximate numeric, then the precision of the result shall be at least as great as the precision of the approximate numeric argument(s).
- e) If <hypothetical set function> is specified, then

Case:

- i) If RANK or DENSE\_RANK is specified, then the declared type of the result is exact numeric with implementation-defined precision and with scale 0 (zero).
  - ii) Otherwise, the declared type of the result is approximate numeric with implementation-defined precision.
- f) If the declared type of the <value expression> simply contained in the <sort specification> of an <inverse distribution function> that specifies PERCENTILE\_CONT is numeric, then the result type is approximate numeric with implementation-defined precision.
  - g) If <array aggregate function> is specified, then the declared type of the result is an array type with implementation-defined maximum cardinality.

56) Subclause 10.10, “<sort specification list>”:

- a) If <null ordering> is not specified, then an implementation-defined <null ordering> is implicit. The implementation-defined default for <null ordering> shall not depend on the context outside of <sort specification list>.

57) Subclause 10.11, “Determination of view privileges”:

- a) If an <authorization identifier> A has UPDATE privilege on every column of a view V, then it is implementation-defined whether A has UPDATE privilege on V.
- b) If an <authorization identifier> A has INSERT privilege on every column of a view V, then it is implementation-defined whether A has INSERT privilege on V.

58) Subclause 11.1, “<schema definition>”:

- a) If <schema character set specification> is not specified, then a <schema character set specification> containing an implementation-defined <character set specification> is implicit.
- b) If <schema path specification> is not specified, then a <schema path specification> containing an implementation-defined <schema name list> is implicit.
- c) If AUTHORIZATION <authorization identifier> is not specified, then an <authorization identifier> equivalent to the implementation-defined <authorization identifier> for the SQL-session is implicit.
- d) The privileges necessary to execute the <schema definition> are implementation-defined.

59) Subclause 11.6, “<table constraint definition>”:

- a) The ordering of the lists of referencing column names and referenced column names in a referential constraint descriptor is implementation-defined, but shall be such that corresponding column names occupy corresponding positions in each list.

60) Subclause 11.34, “<collation definition>”:

- a) The <existing collation name>s that are supported are implementation-defined.
- b) The collation resulting from the specification of EXTERNAL in a <collation definition> may be implementation-defined.

61) Subclause 11.36, “<transliteration definition>”:

- a) The <existing transliteration name>s that are supported are implementation-defined.

62) Subclause 11.40, “<trigger definition>”:

- a) It is implementation-defined whether the <triggered SQL statement> shall not generally contain an <SQL transaction statement>, an <SQL connection statement>, an <SQL schema statement>, an <SQL dynamic statement>, or an <SQL session statement>.
- b) It is implementation-defined whether the <triggered action> shall not generally contain an <SQL data change statement> or a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.

63) Subclause 11.51, “<SQL-invoked routine>”:

- a) If READS SQL DATA is specified, then it is implementation-defined whether the SQL routine body shall not contain an <SQL procedure statement> S that satisfies at least one of the following:
  - i) S is an <SQL data change statement>.



- ii) *S* contains a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data.
  - iii) *S* contains an <SQL procedure statement> that is an <SQL data change statement>.
- b) If CONTAINS SQL is specified, then it is implementation-defined whether the SQL routine body shall not contain an <SQL procedure statement> *S* that satisfies at least one of the following:
  - i) *S* is an <SQL data statement> other than <free locator statement> and <hold locator statement>.
  - ii) *S* contains a <routine invocation> whose subject routine is an SQL-invoked routine that possibly modifies SQL-data or possibly reads SQL-data.
  - iii) *S* contains an <SQL procedure statement> that is an <SQL data statement> other than <free locator statement> and <hold locator statement>.
- c) If DETERMINISTIC is specified, then it is implementation-defined whether the <SQL routine body> shall not contain an <SQL procedure statement> that is possibly non-deterministic.
- d) It is implementation-defined whether the <SQL routine body> shall not contain an <SQL connection statement>, an <SQL schema statement>, an <SQL dynamic statement>, or an <SQL transaction statement> other than a <savepoint statement>, <release savepoint statement>, or a <rollback statement> that specifies a <savepoint clause>.

64) Subclause 11.63, “<sequence generator definition>”:

- a) If <sequence generator data type option> is not specified, then an implementation-defined exact numeric type *DT* with scale 0 (zero) is implicit.

65) Subclause 12.4, “<role definition>”:

- a) The Access Rules are implementation-defined.

66) Subclause 12.7, “<revoke statement>”:

- a) When loss of the USAGE privilege on a character set causes an SQL-client module to be determined to be a lost module, the impact on that SQL-client module is implementation-defined.

67) Subclause 13.1, “<SQL-client module definition>”:

- a) If the explicit or implicit <schema name> does not specify a <catalog name>, then an implementation-defined <catalog name> is implicit.
- b) If <module path specification> is not specified, then a <module path specification> containing an implementation-defined <schema name list> is implicit.
- c) If a <module character set specification> is not specified, then a <module character set specification> that specifies the implementation-defined character set that contains every character that is in <SQL language character> is implicit.

68) Subclause 13.6, “Data type correspondences”:

- a) The C data type that corresponds to an SQL character string type is implementation-defined.
- b) The KIND attribute of the Fortran type CHARACTER that corresponds to an SQL character string type is implementation-defined.

69) Subclause 14.4, “<open statement>”:



- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.

70) Subclause 14.9, “<delete statement: searched>”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.

71) Subclause 14.10, “<truncate table statement>”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.

72) Subclause 14.11, “<insert statement>”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.

73) Subclause 14.12, “<merge statement>”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.

74) Subclause 14.14, “<update statement: searched>”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.

75) Subclause 15.1, “Effect of opening a cursor”:

- a) Whether an SQL-implementation is able to disallow significant changes that would not be visible through a currently open cursor is implementation-defined.

76) Subclause 15.2, “Effect of receiving a result set”:

- a) The operational returnability of a received cursor is implementation-defined.

77) Subclause 15.5, “Effect of a positioned delete”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.
- c) Whether the current row is removed from the sequence of rows of the result set descriptor of *CR* is implementation-defined.

78) Subclause 15.6, “Effect of a positioned update”:

- a) The extent to which an SQL-implementation may disallow independent changes that are not significant is implementation-defined.
- b) If there is any sensitive cursor *CR* that is open in the SQL-transaction in which this statement is being executed and *CR* has been held into a subsequent SQL-transaction, then whether the change resulting from the successful execution of this statement is made visible to *CR* is implementation-defined.

79) Subclause 17.2, “<set transaction statement>”:

- a) The isolation level that is set for a transaction is an implementation-defined isolation level that will not exhibit any of the phenomena that the explicit or implicit <level of isolation> would not exhibit, as specified in Table 8, “SQL-transaction isolation levels and the three phenomena”.

80) Subclause 17.5, “<savepoint statement>”:

- a) The maximum number of savepoints per SQL-transaction is implementation-defined.

81) Subclause 17.7, “<commit statement>”:

- a) Whether a prepared statement, other than the prepared statement of a held cursor, remains in existence is implementation-defined.

82) Subclause 17.8, “<rollback statement>”:

- a) The status of any open cursors in any SQL-client module associated with the current SQL-transaction that were opened by that SQL-transaction before the establishment of a savepoint to which a rollback is executed is implementation-defined.
- b) It is implementation-defined whether or not any prepared statement is deallocated.

83) Subclause 18.1, “<connect statement>”:

- a) If <connection user name> is not specified, then an implementation-defined <connection user name> for the SQL-connection is implicit.
- b) The initial value of the current role name is the null value.
- c) The restrictions on whether or not the <connection user name> shall be identical to the <module authorization identifier> for the SQL-client module that contains the <externally-invoked procedure> that contains the <connect statement> are implementation-defined.
- d) If DEFAULT is specified, then the method by which the default SQL-server is determined is implementation-defined.
- e) The method by which <SQL-server name> is used to determine the appropriate SQL-server is implementation-defined.

## 84) Subclause 19.2, “&lt;set session user identifier statement&gt;”:

- a) Implementation-defined restrictions on the user identifiers to which the SQL-session user identifier can be set are permitted, other than to disallow setting it to its current value.

## 85) Subclause 19.10, “&lt;set session collation statement&gt;”:

- a) If no <character set specification> is specified in a <set session collation statement>, then the character sets for which the SQL-session collations are set are implementation-defined.

## 86) Subclause 20.2, “&lt;allocate descriptor statement&gt;”:

- a) If WITH MAX <occurrences> is not specified, then an implementation-defined default value for <occurrences> that is greater than 0 (zero) is implicit.
- b) The maximum number of SQL descriptor areas and the maximum number of item descriptor areas for a single SQL descriptor area are implementation-defined.

## 87) Subclause 20.5, “&lt;set descriptor statement&gt;”:

- a) Restrictions on changing TYPE, LENGTH, OCTET\_LENGTH, SCALE, COLLATION\_CATALOG, COLLATION\_SCHEMA, COLLATION\_NAME, CHARACTER\_SET\_CATALOG, CHARACTER\_SET\_SCHEMA, and CHARACTER\_SET\_NAME values resulting from the execution of a <describe statement> before execution of an <execute statement>, <dynamic open statement>, or <dynamic fetch statement> are implementation-defined.

## 88) Subclause 20.6, “&lt;prepare statement&gt;”:

- a) The Format and Syntax Rules for a <preparable implementation-defined statement> are implementation-defined.

## 89) Subclause 20.9, “&lt;describe statement&gt;”:

- a) The character set of the data type of <descriptor name> is implementation-defined.
- b) If *SR* does not contain a single routine SQL-invoked *R*, then the values of PARAMETER\_MODE, PARAMETER\_ORDINAL\_POSITION, PARAMETER\_SPECIFIC\_CATALOG, PARAMETER\_SPECIFIC\_SCHEMA, and PARAMETER\_SPECIFIC\_NAME in the descriptor for each <dynamic parameter specification> simply contained in the <call statement> are set to implementation-defined values.

## 90) Subclause 20.11, “&lt;output using clause&gt;”:

- a) If the <output using clause> is used in a <dynamic fetch statement>, then let *CR* be the dynamic cursor identified by the <dynamic fetch statement>, and let *LTDT* be the most specific type of the *i*-th <target specification> or <into argument> of the most recently executed <dynamic fetch statement> prior to the current execution, if any, for the cursor *CR*. It is implementation-defined whether or not an exception condition is raised: *dynamic SQL error — restricted data type attribute violation* if any of several things are true.

## 91) Subclause 21.1, “&lt;embedded SQL host program&gt;”:

- a) If *H* does not contain an <embedded authorization declaration> that specifies SCHEMA, then the <schema name> of the <module authorization clause> of *m* is implementation-defined.
- b) If *H* does not contain an <embedded authorization declaration>, then *m* contains a <module authorization clause> that specifies “SCHEMA *SN*”, where *SN* is an implementation-defined <schema name>.

- c) If an <embedded character set declaration> is not specified, then an <embedded character set declaration> containing an implementation-defined <character set specification> is implicit.
- d) Each <allocate cursor statement> is replaced with a host language procedure or subroutine call of an implementation-defined procedure that associates the <dynamic cursor name> with the prepared statement.
- e) If an <embedded SQL host program> does not contain an <embedded path specification>, then the implied module contains an implementation-defined <module path specification>.

92) Subclause 21.3, “<embedded SQL Ada program>”:

- a) The default character set is implementation-defined.

93) Subclause 21.4, “<embedded SQL C program>”:

- a) The implicit character set in a <C character variable>, a <C VARCHAR variable>, or a <C CLOB variable> is implementation-defined.
- b) The default character set is implementation-defined.

94) Subclause 21.5, “<embedded SQL COBOL program>”:

- a) The COBOL data description clauses, in addition to the PICTURE, SIGN, USAGE and VALUE clauses, that may appear in a <COBOL variable definition> are implementation-defined.
- b) The default character set is implementation-defined.

95) Subclause 21.6, “<embedded SQL Fortran program>”:

- a) The default character set is implementation-defined.

96) Subclause 21.8, “<embedded SQL Pascal program>”:

- a) The default character set is implementation-defined.

97) Subclause 21.9, “<embedded SQL PL/I program>”:

- a) The PL/I data description clauses, in addition to the <PL/I type specification> and the INITIAL clause, that may appear in a <PL/I variable definition> are implementation-defined.
- b) The default character set is implementation-defined.

98) Subclause 22.1, “<direct SQL statement>”:

- a) The <value specification> that represents the null value is implementation-defined.
- b) The Format, Syntax Rules, and Access Rules for <direct implementation-defined statement> are implementation-defined.
- c) Whether a <direct implementation-defined statement> may be associated with an active transaction is implementation-defined.
- d) Whether a <direct implementation-defined statement> initiates a transaction is implementation-defined.

99) Subclause 23.1, “<get diagnostics statement>”:

- a) The maximum length of each variable-length character item in the diagnostics area is in most cases implementation-defined.

- b) The character string value set for CLASS\_ORIGIN and SUBCLASS\_ORIGIN for an implementation-defined class code or subclass code is implementation-defined, but shall not be 'ISO 9075'.
- c) The value of MESSAGE\_TEXT is an implementation-defined character string.
- d) Negative values of COMMAND\_FUNCTION\_CODE are implementation-defined and indicate implementation-defined SQL-statements.
- e) The order in which the occupied condition area numbers are presented when ALL is specified without a condition number is implementation-defined.

100) Subclause 24.1, “SQLSTATE”:

- a) The character set associated with the class value and subclass value of the SQLSTATE parameter is implementation-defined.
- b) The values and meanings for classes and subclasses that begin with one of the <digit>s '5', '6', '7', '8', or '9' or one of the <simple Latin upper case letter>s 'T', 'J', 'K', 'L', 'M', 'N', 'O', 'P', 'Q', 'R', 'S', 'T', 'U', 'V', 'W', 'X', 'Y', or 'Z' are implementation-defined. The values and meanings for all subclasses that are associated with implementation-defined class values are implementation-defined.

101) Clause 25, “Conformance”:

- a) The method of flagging nonconforming SQL language or processing of conforming SQL language is implementation-defined, as is the list of additional <key word>s that may be required by the SQL-implementation.

*(Blank page)*

## Annex C

### (informative)

### Implementation-dependent elements

*This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-3.  
 This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-4.  
 This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-9.  
 This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-10.  
 This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-11.  
 This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-13.  
 This Annex is modified by Annex C, “Implementation-dependent elements”, in ISO/IEC 9075-14.*

This Annex references those places where this part of ISO/IEC 9075 states explicitly that the actions of a conforming SQL-implementation are implementation-dependent.

- 1) Subclause 4.1, “Data types”:
  - a) The physical representation of a value of a data type is implementation-dependent.
  - b) The null value or values for each data type is implementation-dependent.
- 2) Subclause 4.14, “Tables”:
  - a) Because global temporary table contents are distinct within SQL-sessions, and created local temporary tables are distinct within SQL-client modules within SQL-sessions, the effective <schema name> of the schema in which the global temporary table or the created local temporary table is instantiated is an implementation-dependent <schema name> that may be thought of as having been effectively derived from the <schema name> of the schema in which the global temporary table or created local temporary table is defined and the implementation-dependent SQL-session identifier associated with the SQL-session.
  - b) The effective <schema name> of the schema in which the created local temporary table is instantiated may be thought of as being further qualified by a unique implementation-dependent name associated with the SQL-client module in which the created local temporary table is referenced.
- 3) Subclause 4.14.9, “Windowed tables”:
  - a) The window name of a window defined implicitly by an <in-line window specification> is implementation-dependent.
- 4) Subclause 4.27.5, “Result sets returned by SQL-invoked procedures”:
  - a) It is implementation-dependent whether or not result sets whose positions are greater than that maximum number are returned.
- 5) Subclause 4.30, “Diagnostics area”:
  - a) The condition area limit is implementation-dependent when not explicitly specified.

- b) The ordering of the information about conditions placed into the diagnostics area is implementation-dependent, except that the first condition in the diagnostics area always corresponds to the condition corresponding to the SQLSTATE value.
  - c) The maximum number of diagnostics area in a diagnostics area stack is implementation-defined.
- 6) Subclause 4.32, “Cursors”:
- a) If the <cursor specification> does not contain an <order by clause>, or contains an <order by clause> that does not specify the order of the rows completely, then the rows of the result set have an order that is defined only to the extent that the <order by clause> specifies an order and is otherwise implementation-dependent.
  - b) The effect on the position and state of an open cursor when an error occurs during the execution of an SQL-statement that identifies the cursor is implementation-dependent.
  - c) If an asensitive cursor is open and a change is made to SQL-data from within the same SQL-transaction other than through that cursor, then whether that change will be visible through that cursor before it is closed is implementation-dependent.
- 7) Subclause 4.35, “SQL-transactions”:
- a) The schema definitions that are implicitly read on behalf of executing an SQL-statement are implementation-dependent.
- 8) Subclause 4.37, “SQL-sessions”:
- a) A unique implementation-dependent SQL-session identifier is associated with each SQL-session.
- 9) Subclause 4.39, “Client-server operation”:
- a) The <SQL-client module name> of the SQL-client module that is effectively materialized on an SQL-server is implementation-dependent.
  - b) Following the execution of an <SQL procedure statement> by an SQL-server, diagnostic information is passed in an implementation-dependent manner into the SQL-agent's diagnostics area stack in the SQL-client.
  - c) The effect on diagnostic information of incompatibilities between the character repertoires supported by the SQL-client and SQL-server environments is implementation-dependent.
- 10) Subclause 6.7, “<column reference>”:
- a) If *QCR* is a group-invariant column reference and the most specific type of *QCR* is character string, datetime with time zone, or a user-defined type, then *QCR* denotes an implementation-dependent value that is not distinct from the value of *c* in every row of a given group of the qualifying query of *QCR*.
- 11) Subclause 6.9, “<set function specification>”:
- a) The maximum or minimum of a set of values of a user-defined type is implementation-dependent if the comparison of at least two values of the set results in *Unknown*.
- 12) Subclause 6.12, “<cast specification>”:
- a) When a multiset is cast to an array type, the order of elements in the result is implementation-dependent.
- 13) Subclause 6.27, “<numeric value function>”:



- a) If <regex occurrences function> specifies OCTETS and the octet whose position is given by the <start position> is not the first octet of a character, then the result of <regex occurrences function> is implementation-dependent.
- b) If <regex position expression> specifies OCTETS and the octet whose position is given by the <start position> is not the first octet of a character, then the result of <regex position expression> is implementation-dependent.

14) Subclause 6.29, “<string value function>”:

- a) If <regex substring function> specifies OCTETS and the octet whose position is given by the <start position> is not the first octet of a character, then the result of <regex occurrences function> is implementation-dependent.
- b) If <regex transliteration> specifies OCTETS and the octet whose position is given by the <start position> is not the first octet of a character, then the result of <regex position expression> is implementation-dependent.

15) Subclause 6.31, “<datetime value function>”:

- a) The time of evaluation of the CURRENT\_DATE, CURRENT\_TIME, and CURRENT\_TIMESTAMP functions during the execution of an SQL-statement is implementation-dependent.

16) Subclause 6.32, “<interval value expression>”:

- a) The start datetime used for converting intervals to scalars for subtraction purposes is implementation-dependent.

17) Subclause 6.36, “<array value constructor>”:

- a) The order of array elements in the result of an <array value constructor by query> which is not decided by the <order by clause> is implementation-dependent.

18) Subclause 7.1, “<row value constructor>”:

- a) The names of the fields of a <row value constructor> that specifies a <row value constructor element list> are implementation-dependent.
- b) The names of the fields of a <contextually typed row value constructor> are implementation-dependent.

19) Subclause 7.3, “<table value constructor>”:

- a) The column names of a <table value constructor> or a <contextually typed table value constructor> are implementation-dependent.

20) Subclause 7.9, “<group by clause>”:

- a) If the declared type of a grouping column is a user-defined type and the comparison of that column for two rows results in *Unknown*, then the assignment of those rows to groups in the result of the <group by clause> is implementation-dependent.
- b) When a <search condition> or <value expression> is applied to a group, the value of a group-invariant column reference whose most specific type is character string, datetime with time zone or a user-defined type, and that references a column that is functionally dependent on the grouping columns is implementation-dependent.

21) Subclause 7.11, “<window clause>”:

- a) If the window ordering clause of a window structure descriptor is absent, then the window ordering is implementation-dependent.
- b) The window ordering of peer rows within a window partition is implementation-dependent, but the window ordering shall be the same for all window structure descriptors that are order-equivalent. It shall also be the same for windows *W1* and *W2* if *W1* is the ordering window for *W2*.

22) Subclause 7.12, “<query specification>”:

- a) When a column is not named by an <as clause> and is not derived from a single column reference, then the name of the column is implementation-dependent.
- b) If the <set quantifier> DISTINCT is specified, and the most specific type of a result column is character string, datetime with time zone or a user-defined type, then the precise values retained in that column after eliminating redundant duplicates is implementation-dependent.

23) Subclause 7.13, “<query expression>”:

- a) If a <simple table> is neither a <query specification> nor an <explicit table>, then the name of each column of the <simple table> is implementation-dependent.
- b) If a <query term> immediately contains INTERSECT and the <column name>s of a pair of corresponding columns of the operand tables are not equivalent, then the result column has an implementation-dependent <column name>.
- c) If a <query expression body> immediately contains UNION or INTERSECT, and the <column name>s of a pair of corresponding columns of the operand tables are not equivalent, then the result column has an implementation-dependent <column name>.
- d) If a <query expression> does not simply contain an <order by clause>, then the ordering of rows in the table specified by that <query expression> is implementation-dependent.
- e) If a <query expression> immediately contains an <order by clause> and a group of two or more rows in the table specified by that <query expression> contain values that are not distinct in all sort keys specified in the <order by clause>, then the ordering of these rows in that group is implementation-dependent.
- f) The relative ordering of two non-null values of a user-defined type *UDT* whose comparison as determined by the user-defined ordering of *UDT* is *Unknown* is implementation-dependent.

24) Subclause 8.2, “<comparison predicate>”:

- a) When the operations MAX, MIN, DISTINCT, and references to a grouping column refer to a variable-length character string, the specific value selected from the set of equal values is implementation-dependent.

- 25) a) If the result of comparing *V1* and *V2* for equality according to Subclause 8.2, “<comparison predicate>”, is *Unknown*, then it is implementation-dependent whether the result is *True* or *False*.

26) Subclause 10.4, “<routine invocation>”:

- a) Each SQL argument *A<sub>i</sub>* in *SAL* is evaluated, in an implementation-dependent order, to obtain a value *V<sub>i</sub>*.
- b) If *P* is not a conforming program according to the standard or other specification for the language of *R*, then the results of any execution of *P* are implementation-dependent.

## 27) Subclause 10.9, “&lt;aggregate function&gt;”:

- a) If the declared type of the argument of MAX or MIN is a user-defined type and the comparison of two values results in Unknown, then the maximum or minimum is implementation-dependent.
- b) If an <array aggregate function> does not specify ORDER BY, then the ordering of the elements in the result array is implementation-dependent.

## 28) Subclause 10.10, “&lt;sort specification list&gt;”:

- a) If  $PV_i$  and  $QV_i$  are not the null value and the result of “ $PV_i$  <comp op>  $QV_i$ ” is Unknown, then the relative ordering of  $PV_i$  and  $QV_i$  is implementation-dependent.
- b) The relative ordering of two rows that are not distinct with respect to the <sort specification> is implementation-dependent.

## 29) Subclause 11.6, “&lt;table constraint definition&gt;”:

- a) The <constraint name> of a constraint that does not specify a <constraint name definition> is implementation-dependent.

## 30) Subclause 11.8, “&lt;referential constraint definition&gt;”:

- a) The specific value to use for cascading among various values that are not distinct is implementation-dependent.

## 31) Subclause 11.25, “&lt;domain definition&gt;”:

- a) The <constraint name> of a constraint that does not specify a <constraint name definition> is implementation-dependent.

## 32) Subclause 11.34, “&lt;collation definition&gt;”:

- a) The collation of characters for which a collation is not otherwise specified is implementation-dependent.

## 33) Subclause 14.5, “&lt;fetch statement&gt;”:

- a) The order of assignment to targets in the <fetch target list> of values returned by a <fetch statement>, other than status parameters, is implementation-dependent.
- b) If an error occurs during assignment of a value to a target during the execution of a <select statement: single row>, then the values of targets other than status parameters are implementation-dependent.
- c) If an exception condition occurs during the assignment of a value to a target, then the values of all targets are implementation-dependent and *CR* remains positioned on the current row.
- d) It is implementation-dependent whether *CR* remains positioned on the current row when an exception condition is raised during the derivation of any <derived column>.

## 34) Subclause 14.7, “&lt;select statement: single row&gt;”:

- a) The order of assignment to targets in the <select target list> of values returned by a <select statement: single row>, other than status parameters, is implementation-dependent.
- b) If the cardinality of the <query expression> is greater than 1 (one), then it is implementation-dependent whether or not values are assigned to the targets identified by the <select target list>.

- c) If an error occurs during assignment of a value to a target during the execution of a <select statement: single row>, then the values of targets other than status parameters are implementation-dependent.

35) Subclause 14.11, “<insert statement>”:

- a) The generation of the value of a derived self-referencing column is implementation-dependent.

36) Subclause 14.12, “<merge statement>”:

- a) The generation of the value of a derived self-referencing column is implementation-dependent.

37) Subclause 17.2, “<set transaction statement>”:

- a) If <number of conditions> is not specified, then an implementation-dependent value not less than 1 (one) is implicit.
- b) If <number of conditions> is specified and is greater than the implementation-dependent maximum value for <number of conditions>, then that maximum value is used.

38) Subclause 18.3, “<disconnect statement>”:

- a) If ALL is specified, then *l* is a list representing every active SQL-connection that has been established by a <connect statement> by the current SQL-agent and that has not yet been disconnected by a <disconnect statement>, in an implementation-dependent order.

39) Subclause 20.4, “<get descriptor statement>”:

- a) If an exception condition is raised in a <get descriptor statement>, then the values of all targets specified by <simple target specification 1> and <simple target specification 2> are implementation-dependent.
- b) For a <dynamic parameter specification>, the value of UNNAMED is 1 (one) and the value of NAME is implementation-dependent.
- c) The value retrieved by a <get descriptor statement> for any field whose value is undefined is implementation-dependent.

40) Subclause 20.5, “<set descriptor statement>”:

- a) If an exception condition is raised in a <set descriptor statement>, then the values of all elements of the descriptor specified in the <set descriptor statement> are implementation-dependent.

41) Subclause 20.10, “<input using clause>”:

- a) When a <describe input statement> is used, the values for NAME, DATA, and INDICATOR in the SQL dynamic descriptor area structure is implementation-dependent. If TYPE indicates a character string type or a binary large object type, then the values of SCALE and PRECISION are implementation-dependent. If TYPE indicates an exact or approximate numeric type, then the values of LENGTH and OCTET\_LENGTH are implementation-dependent. If TYPE indicates a boolean type, then the values of PRECISION, SCALE, LENGTH, and OCTET\_LENGTH are implementation-dependent.

42) Subclause 20.11, “<output using clause>”:

- a) When a <describe output statement> is executed, the values of DATA and INDICATOR are implementation-dependent. If TYPE indicates a character string type or a binary large object type, then the values of SCALE and PRECISION are implementation-dependent. If TYPE indicates an exact or approximate numeric type, then the values of LENGTH and OCTET\_LENGTH are implementation-dependent. If TYPE indicates a boolean type, then the values of PRECISION, SCALE, LENGTH, and OCTET\_LENGTH are implementation-dependent.

## 43) Subclause 20.13, “&lt;execute immediate statement&gt;”:

- a) The <statement name> of the statement prepared when an <execute immediate statement> is executed is implementation-dependent.

## 44) Subclause 21.1, “&lt;embedded SQL host program&gt;”:

- a) The <SQL-client module name> of the implied <SQL-client module definition> derived from an <embedded SQL host program> is implementation-dependent.
- b) If an <embedded SQL host program> does not contain an <embedded authorization declaration>, then the <module authorization identifier> of the implied <SQL-client module definition> derived from the <embedded SQL host program> is implementation-dependent.
- c) In each <declare cursor> in the implied <SQL-client module definition> derived from an <embedded SQL host program>, each <embedded variable name> has been replaced consistently with a distinct <host parameter name> that is implementation-dependent.
- d) The <procedure name> of each <externally-invoked procedure> in the implied <SQL-client module definition> derived from an <embedded SQL host program> is implementation-dependent.
- e) In each <externally-invoked procedure> in the implied <SQL-client module definition> derived from an <embedded SQL host program>, each <embedded variable name> has been replaced consistently with a distinct <host parameter name> that is implementation-dependent.
- f) For <SQL procedure statement>s other than <open statement>s, whether one <externally-invoked procedure> in the implied <SQL-client module definition> derived from an <embedded SQL host program> can correspond to more than one <SQL procedure statement> in the <embedded SQL host program> is implementation-dependent.
- g) In each <externally-invoked procedure> in the implied <SQL-client module definition> derived from an <embedded SQL host program>, the order of the instances of <host parameter declaration> is implementation-dependent.

## 45) Subclause 22.1, “&lt;direct SQL statement&gt;”:

- a) A <commit statement> or a <rollback statement> is executed. If an unrecoverable error has occurred, or if the direct invocation of SQL terminated unexpectedly, or if any constraint is not satisfied, then a <rollback statement> is performed. Otherwise, the choice of which of these SQL-statements to perform is implementation-dependent. The determination of whether a direct invocation of SQL has terminated unexpectedly is implementation-dependent.

## 46) Subclause 23.1, “&lt;get diagnostics statement&gt;”:

- a) The value of ROW\_COUNT following the execution of an SQL-statement that does not directly result in the execution of a <delete statement: searched>, an <insert statement>, a <merge statement>, or an <update statement: searched> is implementation-dependent.
- b) If <condition number> has a value other than 1 (one), then the association between <condition number> values and specific conditions raised during evaluation of the General Rules for that SQL-statement is implementation-dependent.
- c) The order in which the statement information items and the condition information items of a particular condition area are presented when ALL is specified is implementation-dependent.

*(Blank page)*

## **Annex D** (informative)

### **Deprecated features**

*This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-1.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-3.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-4.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-9.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-10.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-11.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-13.  
This Annex is modified by Annex D, “Deprecated features”, in ISO/IEC 9075-14.*

It is intended that the following features will be removed at a later date from a revised version of this part of ISO/IEC 9075:

*None.*

*(Blank page)*



## Annex E

### (informative)

### Incompatibilities with ISO/IEC 9075:2003

*This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075:2003”, in ISO/IEC 9075-3.  
 This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075:2003”, in ISO/IEC 9075-4.  
 This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075:2003”, in ISO/IEC 9075-9.  
 This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075:2003”, in ISO/IEC 9075-10.  
 This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075:2003”, in ISO/IEC 9075-11.  
 This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075:2003”, in ISO/IEC 9075-13.  
 This Annex is modified by Annex E, “Incompatibilities with ISO/IEC 9075-14:2006”, in ISO/IEC 9075-14.*

This edition of this part of ISO/IEC 9075 introduces some incompatibilities with the earlier version of Database Language SQL as specified in ISO/IEC 9075-2:2003.

Except as specified in this Annex, features and capabilities of Database Language SQL are compatible with ISO/IEC 9075-2:2003.

- 1) In ISO/IEC 9075-2:2003, <search or cycle clause>'s <recursive search order> required the specification of <sort specification list>. This edition of ISO/IEC 9075-2 further restricts that specification to <column name list>.
- 2) In ISO/IEC 9075-2:2003, the declared type of a <binary string literal> was binary large object string type. In this edition of ISO/IEC 9075-2, it is implementation-defined whether the declared type of a <binary string literal> is binary large object string type, variable-length binary string, or fixed-length binary string.
- 3) In ISO/IEC 9075-2:2003, the <non-reserved word> EXCEPTION could be used as a synonym for the <non-reserved word> CONDITION in a <get diagnostics statement>. In this edition of ISO/IEC 9075-2, the <non-reserved word> EXCEPTION has been omitted entirely.
- 4) In ISO/IEC 9075-2:2003, the exception raised by <dynamic open statement> when there is an attempt to open a declared dynamic cursor for a prepared statement that is not a <cursor specification> was *dynamic SQL error — invalid SQL statement name*. This exception could not actually occur, owing to a General Rule for <prepare statement>. In this edition of ISO/IEC 9075-2, the exception can occur, but has been changed to *dynamic SQL error — prepared statement not a cursor specification*.
- 5) A number of additional <reserved word>s have been added to the language. These <reserved word>s are:
  - ARRAY\_AGG
  - CURRENT\_CATALOG
  - CURRENT\_SCHEMA
  - LIKE\_REGEX
  - OCCURRENCES\_REGEX
  - POSITION\_REGEX

- SUBSTRING\_REGEX
- TRANSLATE\_REGEX
- TRUNCATE
- VARBINARY

## Annex F

### (informative)

### SQL feature taxonomy

*This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-3.  
 This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-4.  
 This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-9.  
 This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-10.  
 This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-11.  
 This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-13.  
 This Annex is modified by Annex F, “SQL feature taxonomy”, in ISO/IEC 9075-14.*

This Annex describes a taxonomy of features and packages defined in this part of ISO/IEC 9075.

Table 36, “Feature taxonomy and definition for mandatory features”, contains a taxonomy of the mandatory features of SQL language that are specified in this part of ISO/IEC 9075. Table 37, “Feature taxonomy for optional features”, contains a taxonomy of the optional features of the SQL language that are specified in this part of ISO/IEC 9075.

In these tables, the first column contains a counter that may be used to quickly locate rows of the table; these values otherwise have no use and are not stable — that is, they are subject to change in future editions of or even Technical Corrigenda to ISO/IEC 9075 without notice.

The “Feature ID” column of Table 36, “Feature taxonomy and definition for mandatory features”, and of Table 37, “Feature taxonomy for optional features”, specifies the formal identification of each feature and each subfeature contained in the table.

The “Feature Name” column of Table 36, “Feature taxonomy and definition for mandatory features”, contains a brief description of the feature or subfeature associated with the Feature ID value.

The “Feature Description” column of Table 36, “Feature taxonomy and definition for mandatory features”, provides the only definition of the mandatory features of this part of ISO/IEC 9075. This definition consists of indications of specific language elements supported in each feature, subject to the constraints of all Syntax Rules, Access Rules, and Conformance Rules.

**Table 36 — Feature taxonomy and definition for mandatory features**

	Feature ID	Feature Name	Feature Description
1	E011	Numeric data types	Subclause 6.1, “<data type>”, <numeric type>, including numeric expressions, numeric literals, numeric comparisons, and numeric assignments

	Feature ID	Feature Name	Feature Description
2	E011-01	INTEGER and SMALLINT data types (including all spellings)	— Subclause 5.2, “<token> and <separator>”: The <reserved word>s INT, INTEGER, and SMALLINT — Subclause 5.3, “<literal>”: [<sign>] <unsigned integer> — Subclause 6.1, “<data type>”: The INTEGER and SMALLINT <exact numeric type>s — Subclause 13.6, “Data type correspondences”: Type correspondences for INTEGER and SMALLINT for all supported languages
3	E011-02	REAL, DOUBLE PRECISION, and FLOAT data types	— Subclause 5.2, “<token> and <separator>”: The <reserved word>s REAL, DOUBLE, FLOAT, and PRECISION — Subclause 5.3, “<literal>”: [<sign>] <approximate numeric literal> — Subclause 6.1, “<data type>”: <approximate numeric type> — Subclause 13.6, “Data type correspondences”: Type correspondences for REAL, DOUBLE PRECISION, and FLOAT for all supported languages
4	E011-03	DECIMAL and NUMERIC data types	— Subclause 5.2, “<token> and <separator>”: The <reserved word>s DEC, DECIMAL, and NUMERIC — Subclause 5.3, “<literal>”: [<sign>] <exact numeric literal> — Subclause 6.1, “<data type>”: The DECIMAL and NUMERIC <exact numeric type>s — Subclause 13.6, “Data type correspondences”: Type correspondences for DECIMAL and NUMERIC for all supported languages
5	E011-04	Arithmetic operators	— Subclause 6.26, “<numeric value expression>”: When the <numeric primary> is a <value expression primary>
6	E011-05	Numeric comparison	— Subclause 8.2, “<comparison predicate>”: For the numeric data types, without support for <table subquery> and without support for Feature F131, “Grouped operations”
7	E011-06	Implicit casting among the numeric data types	— Subclause 8.2, “<comparison predicate>”: Values of any of the numeric data types can be compared to each other; such values are compared with respect to their algebraic values — Subclause 9.1, “Retrieval assignment”, and Subclause 9.2, “Store assignment”: Values of one numeric type can be assigned to another numeric type, subject to rounding, truncation, and out of range conditions

	Feature ID	Feature Name	Feature Description
8	E021	Character string types	— Subclause 6.1, “<data type>”: <character string type>, including character expressions, character literals, character comparisons, character assignments, and other operations on character data
9	E021-01	CHARACTER data type (including all its spellings)	— Subclause 5.2, “<token> and <separator>”: The <reserved word>s CHAR and CHARACTER — Subclause 6.1, “<data type>”: The CHARACTER <character string type> — Subclause 6.28, “<string value expression>”: For values of type CHARACTER — Subclause 13.6, “Data type correspondences”: Type correspondences for CHARACTER for all supported languages
10	E021-02	CHARACTER VARYING data type (including all its spellings)	— Subclause 5.2, “<token> and <separator>”: The <reserved word>s VARCHAR and VARYING — Subclause 6.1, “<data type>”: The CHARACTER VARYING <character string type> — Subclause 6.28, “<string value expression>”: For values of type CHARACTER VARYING — Subclause 13.6, “Data type correspondences”: Type correspondences for CHARACTER VARYING for all supported languages
11	E021-03	Character literals	— Subclause 5.3, “<literal>”: <quote> [ <character representation>... ] <quote>
12	E021-04	CHARACTER_LENGTH function	— Subclause 6.27, “<numeric value function>”: The <char length expression>
13	E021-05	OCTET_LENGTH function	— Subclause 6.27, “<numeric value function>”: The <octet length expression>
14	E021-06	SUBSTRING function	— Subclause 6.29, “<string value function>”: The <character substring function>
15	E021-07	Character concatenation	— Subclause 6.28, “<string value expression>”: The <concatenation> expression
16	E021-08	UPPER and LOWER functions	— Subclause 6.29, “<string value function>”: The <fold> function
17	E021-09	TRIM function	— Subclause 6.29, “<string value function>”: The <trim function>

	Feature ID	Feature Name	Feature Description
18	E021-10	Implicit casting among the fixed-length and variable-length character string types	— Subclause 8.2, “<comparison predicate>”: Values of either the CHARACTER or CHARACTER VARYING data types can be compared to each other — Subclause 9.1, “Retrieval assignment”, and Subclause 9.2, “Store assignment”: Values of either the CHARACTER or CHARACTER VARYING data type can be assigned to the other type, subject to truncation conditions
19	E021-11	POSITION function	— Subclause 6.27, “<numeric value function>”: The <position expression>
20	E021-12	Character comparison	— Subclause 8.2, “<comparison predicate>”: For the CHARACTER and CHARACTER VARYING data types, without support for <table subquery> and without support for Feature F131, “Grouped operations”
21	E031	Identifiers	— Subclause 5.2, “<token> and <separator>”: <regular identifier> and <delimited identifier>
22	E031-01	Delimited identifiers	— Subclause 5.2, “<token> and <separator>”: <delimited identifier>
23	E031-02	Lower case identifiers	— Subclause 5.2, “<token> and <separator>”: An alphabetic character in a <regular identifier> can be either lower case or upper case (meaning that non-delimited identifiers need not comprise only upper case letters)
24	E031-03	Trailing underscore	— Subclause 5.2, “<token> and <separator>”: The list <identifier part> in a <regular identifier> can be an <underscore>
25	E051	Basic query specification	— Subclause 7.12, “<query specification>”: When <table reference> is a <table or query name> that is a <table name>, without the support of Feature F131, “Grouped operations”
26	E051-01	SELECT DISTINCT	— Subclause 7.12, “<query specification>”: With a <set quantifier> of DISTINCT, but without subfeatures E051-02 through E051-09

	Feature ID	Feature Name	Feature Description
27	E051-02	GROUP BY clause	— Subclause 7.4, “<table expression>”: <group by clause>, but without subfeatures E051-04 through E051-09 — Subclause 7.9, “<group by clause>”: With the restrictions that the <group by clause> shall contain all non-aggregated columns in the <select list> and that any column in the <group by clause> shall also appear in the <select list>
28	E051-04	GROUP BY can contain columns not in <select list>	— Subclause 7.9, “<group by clause>”: Without the restriction that any column in the <group by clause> shall also appear in the <select list>
29	E051-05	Select list items can be renamed	— Subclause 7.12, “<query specification>”: <as clause>
30	E051-06	HAVING clause	— Subclause 7.4, “<table expression>”: <having clause> — Subclause 7.10, “<having clause>”
31	E051-07	Qualified * in select list	— Subclause 7.12, “<query specification>”: <qualified asterisk>
32	E051-08	Correlation names in the FROM clause	— Subclause 7.6, “<table reference>”: [ AS ] <correlation name>
33	E051-09	Rename columns in the FROM clause	— Subclause 7.6, “<table reference>”: [ AS ] <correlation name> [ <left paren> <derived column list> <right paren> ]
34	<b>E061</b>	<b>Basic predicates and search conditions</b>	— Subclause 8.20, “<search condition>”, and Subclause 8.1, “<predicate>”
35	E061-01	Comparison predicate	— Subclause 8.2, “<comparison predicate>”: For supported data types, without support for <table subquery>
36	E061-02	BETWEEN predicate	— Subclause 8.3, “<between predicate>”
37	E061-03	IN predicate with list of values	— Subclause 8.4, “<in predicate>”: Without support for <table subquery>
38	E061-04	LIKE predicate	— Subclause 8.5, “<like predicate>”: Without [ ESCAPE <escape character> ]
39	E061-05	LIKE predicate: ESCAPE clause	— Subclause 8.5, “<like predicate>”: With [ ESCAPE <escape character> ]

	Feature ID	Feature Name	Feature Description
40	E061-06	NULL predicate	— Subclause 8.8, “<null predicate>”: Without Feature F481, “Expanded NULL predicate”
41	E061-07	Quantified comparison predicate	— Subclause 8.9, “<quantified comparison predicate>”: Without support for <table subquery>
42	E061-08	EXISTS predicate	— Subclause 8.10, “<exists predicate>”
43	E061-09	Subqueries in comparison predicate	— Subclause 8.2, “<comparison predicate>”: For supported data types, with support for <table subquery>
44	E061-11	Subqueries in IN predicate	— Subclause 8.4, “<in predicate>”: With support for <table subquery>
45	E061-12	Subqueries in quantified comparison predicate	— Subclause 8.9, “<quantified comparison predicate>”: With support for <table subquery>
46	E061-13	Correlated subqueries	— Subclause 8.1, “<predicate>”: When a <correlation name> can be used in a <table subquery> as a correlated reference to a column in the outer query
47	E061-14	Search condition	— Subclause 8.20, “<search condition>”
48	E071	Basic query expressions	— Subclause 7.13, “<query expression>”
49	E071-01	UNION DISTINCT table operator	— Subclause 7.13, “<query expression>”: With support for UNION [ DISTINCT ]
50	E071-02	UNION ALL table operator	— Subclause 7.13, “<query expression>”: With support for UNION ALL
51	E071-03	EXCEPT DISTINCT table operator	— Subclause 7.13, “<query expression>”: With support for EXCEPT [ DISTINCT ]
52	E071-05	Columns combined via table operators need not have exactly the same data type.	— Subclause 7.13, “<query expression>”: Columns combined via UNION and EXCEPT need not have exactly the same data type
53	E071-06	Table operators in subqueries	— Subclause 7.13, “<query expression>”: <table subquery>s can specify UNION and EXCEPT
54	E081	Basic Privileges	— Subclause 12.3, “<privileges>”
55	E081-01	SELECT privilege at the table level	— Subclause 12.3, “<privileges>”: With <action> of SELECT without <privilege column list>



	Feature ID	Feature Name	Feature Description
56	E081-02	DELETE privilege	— Subclause 12.3, “<privileges>”: With <action> of DELETE
57	E081-03	INSERT privilege at the table level	— Subclause 12.3, “<privileges>”: With <action> of INSERT without <privilege column list>
58	E081-04	UPDATE privilege at the table level	— Subclause 12.3, “<privileges>”: With <action> of UPDATE without <privilege column list>
59	E081-05	UPDATE privilege at the column level	— Subclause 12.3, “<privileges>”: With <action> of UPDATE <left paren> <privilege column list> <right paren>
60	E081-06	REFERENCES privilege at the table level	— Subclause 12.3, “<privileges>”: with <action> of REFERENCES without <privilege column list>
61	E081-07	REFERENCES privilege at the column level	— Subclause 12.3, “<privileges>”: With <action> of REFERENCES <left paren> <privilege column list> <right paren>
62	E081-08	WITH GRANT OPTION	— Subclause 12.2, “<grant privilege statement>”: WITH GRANT OPTION
63	E081-09	USAGE privilege	— Subclause 12.3, “<privileges>”: With <action> of USAGE
64	E081-10	EXECUTE privilege	— Subclause 12.3, “<privileges>”: With <action> of EXECUTE
65	E091	Set functions	— Subclause 6.9, “<set function specification>”
66	E091-01	AVG	— Subclause 6.9, “<set function specification>”: With <computational operation> of AVG
67	E091-02	COUNT	— Subclause 6.9, “<set function specification>”: With <computational operation> of COUNT
68	E091-03	MAX	— Subclause 6.9, “<set function specification>”: With <computational operation> of MAX
69	E091-04	MIN	— Subclause 6.9, “<set function specification>”: With <computational operation> of MIN
70	E091-05	SUM	— Subclause 6.9, “<set function specification>”: With <computational operation> of SUM
71	E091-06	ALL quantifier	— Subclause 6.9, “<set function specification>”: With <set quantifier> of ALL

	Feature ID	Feature Name	Feature Description
72	E091-07	DISTINCT quantifier	— Subclause 6.9, “<set function specification>”: With <set quantifier> of DISTINCT
73	E101	Basic data manipulation	— Clause 14, “Data manipulation”: <insert statement>, <delete statement: searched>, and <update statement: searched>
74	E101-01	INSERT statement	— Subclause 14.11, “<insert statement>”: When a <contextually typed table value constructor> can consist of no more than a single <contextually typed row value expression>
75	E101-03	Searched UPDATE statement	— Subclause 14.14, “<update statement: searched>”: But without support either of Feature E153, “Updatable tables with subqueries”, or Feature F221, “Explicit defaults”
76	E101-04	Searched DELETE statement	— Subclause 14.9, “<delete statement: searched>”
77	E111	Single row SELECT statement	— Subclause 14.7, “<select statement: single row>”: Without support of Feature F131, “Grouped operations”
78	E121	Basic cursor support	— Clause 14, “Data manipulation”: <declare cursor>, <open statement>, <fetch statement>, <close statement>, <delete statement: positioned>, and <update statement: positioned>
79	E121-01	DECLARE CURSOR	— Subclause 14.1, “<declare cursor>”: When each <value expression> in the <sort key> shall be a <column reference> and that <column reference> shall also be in the <select list>, and <cursor holdability> is not specified
80	E121-02	ORDER BY columns need not be in select list	— Subclause 14.1, “<declare cursor>”: Extend subfeature E121-01 so that <column reference> need not also be in the <select list>
81	E121-03	Value expressions in ORDER BY clause	— Subclause 14.1, “<declare cursor>”: Extend subfeature E121-01 so that the <value expression> in the <sort key> need not be a <column reference>
82	E121-04	OPEN statement	— Subclause 14.4, “<open statement>”

	Feature ID	Feature Name	Feature Description
83	E121-06	Positioned UPDATE statement	— Subclause 14.13, “<update statement: positioned>”: Without support of either Feature E153, “Updateable tables with subqueries” or Feature F221, “Explicit defaults”
84	E121-07	Positioned DELETE statement	— Subclause 14.8, “<delete statement: positioned>”
85	E121-08	CLOSE statement	— Subclause 14.6, “<close statement>”
86	E121-10	FETCH statement: implicit NEXT	— Subclause 14.5, “<fetch statement>”
87	E121-17	WITH HOLD cursors	— Subclause 14.1, “<declare cursor>”: Where the <value expression> in the <sort key> need not be a <column reference> and need not be in the <select list>, and <cursor holdability> may be specified
88	E131	Null value support (nulls in lieu of values)	— Subclause 4.13, “Columns, fields, and attributes”: Nullability characteristic — Subclause 6.5, “<contextually typed value specification>”: <null specification>
89	E141	Basic integrity constraints	— Subclause 11.6, “<table constraint definition>”: As specified by the subfeatures of this feature in this table
90	E141-01	NOT NULL constraints	— Subclause 11.4, “<column definition>”: With <column constraint> of NOT NULL
91	E141-02	UNIQUE constraints of NOT NULL columns	— Subclause 11.4, “<column definition>”: With <unique specification> of UNIQUE for columns specified as NOT NULL — Subclause 11.7, “<unique constraint definition>”: With <unique specification> of UNIQUE
92	E141-03	PRIMARY KEY constraints	— Subclause 11.4, “<column definition>”: With <unique specification> of PRIMARY KEY for columns specified as NOT NULL — Subclause 11.7, “<unique constraint definition>”: With <unique specification> of PRIMARY KEY

	Feature ID	Feature Name	Feature Description
93	E141-04	Basic FOREIGN KEY constraint with the NO ACTION default for both referential delete action and referential update action.	— Subclause 11.4, “<column definition>”: With <column constraint> of <references specification> — Subclause 11.8, “<referential constraint definition>”: Where the columns in the <column name list>, if specified, shall be in the same order as the names in the <unique column list> of the applicable <unique constraint definition> and the <data type>s of the matching columns shall be the same
94	E141-06	CHECK constraints	— Subclause 11.4, “<column definition>”: With <column constraint> of <check constraint definition> — Subclause 11.9, “<check constraint definition>”
95	E141-07	Column defaults	— Subclause 11.4, “<column definition>”: With <default clause>
96	E141-08	NOT NULL inferred on PRIMARY KEY	— Subclause 11.4, “<column definition>”, and Subclause 11.7, “<unique constraint definition>”: Remove the restriction in subfeatures E141-02 and E141-03 that NOT NULL be specified along with every PRIMARY KEY and UNIQUE constraint — Subclause 11.4, “<column definition>”: NOT NULL is implicit on PRIMARY KEY constraints
97	E141-10	Names in a foreign key can be specified in any order	— Subclause 11.4, “<column definition>”, and Subclause 11.8, “<referential constraint definition>”: Extend subfeature E141-04 so that the columns in the <column name list>, if specified, need not be in the same order as the names in the <unique column list> of the applicable <unique constraint definition>
98	E151	Transaction support	— Clause 17, “Transaction management”: <commit statement> and <rollback statement>
99	E151-01	COMMIT statement	— Subclause 17.7, “<commit statement>”
100	E151-02	ROLLBACK statement	— Subclause 17.8, “<rollback statement>”
101	E152	Basic SET TRANSACTION statement	— Subclause 17.2, “<set transaction statement>”
102	E152-01	SET TRANSACTION statement: ISOLATION LEVEL SERIALIZABLE clause	— Subclause 17.2, “<set transaction statement>”: With <transaction mode> of ISOLATION LEVEL SERIALIZABLE clause

	Feature ID	Feature Name	Feature Description
103	E152-02	SET TRANSACTION statement: READ ONLY and READ WRITE clauses	— Subclause 17.2, “<set transaction statement>”: with <transaction access mode> of READ ONLY or READ WRITE
104	E153	Updatable queries with subqueries	— Subclause 7.13, “<query expression>”: A <query expression> is updatable even though its <where clause> contains a <subquery>
105	E161	SQL comments using leading double minus	— Subclause 5.2, “<token> and <separator>”: <simple comment>
106	E171	SQLSTATE support	— Subclause 24.1, “SQLSTATE”
107	E182	Module language	— Clause 13, “SQL-client modules”  NOTE 509 — An SQL-implementation is required to supply at least one binding to a standard host language using either module language, embedded SQL, or both.
108	F031	Basic schema manipulation	— Clause 11, “Schema definition and manipulation”: Selected facilities as indicated by the sub-features of this Feature
109	F031-01	CREATE TABLE statement to create persistent base tables	— Subclause 11.3, “<table definition>”: Not in the context of a <schema definition>
110	F031-02	CREATE VIEW statement	— Subclause 11.23, “<view definition>”: Not in the context of a <schema definition>, and without support of Feature F081, “UNION and EXCEPT in views”
111	F031-03	GRANT statement	— Subclause 12.1, “<grant statement>”: Not in the context of a <schema definition>
112	F031-04	ALTER TABLE statement: ADD COLUMN clause	— Subclause 11.10, “<alter table statement>”: The <add column definition> clause — Subclause 11.11, “<add column definition>”
113	F031-13	DROP TABLE statement: RESTRICT clause	— Subclause 11.22, “<drop table statement>”: With a <drop behavior> of RESTRICT
114	F031-16	DROP VIEW statement: RESTRICT clause	— Subclause 11.24, “<drop view statement>”: With a <drop behavior> of RESTRICT
115	F031-19	REVOKE statement: RESTRICT clause	— Subclause 12.7, “<revoke statement>”: With a <drop behavior> of RESTRICT, only where the use of this statement can be restricted to the owner of the table being dropped
116	F041	Basic joined table	— Subclause 7.7, “<joined table>”

	Feature ID	Feature Name	Feature Description
117	F041-01	Inner join (but not necessarily the INNER keyword)	— Subclause 7.6, “<table reference>”: The <joined table> clause, but without support for subfeatures F041-02 through F041-08
118	F041-02	INNER keyword	— Subclause 7.7, “<joined table>”: <join type> of INNER
119	F041-03	LEFT OUTER JOIN	— Subclause 7.7, “<joined table>”: <outer join type> of LEFT
120	F041-04	RIGHT OUTER JOIN	— Subclause 7.7, “<joined table>”: <outer join type> of RIGHT
121	F041-05	Outer joins can be nested	— Subclause 7.7, “<joined table>”: Subfeature F041-1 extended so that a <table reference> within the <joined table> can itself be a <joined table>
122	F041-07	The inner table in a left or right outer join can also be used in an inner join	— Subclause 7.7, “<joined table>”: Subfeature F041-1 extended so that a <table name> within a nested <joined table> can be the same as a <table name> in an outer <joined table>
123	F041-08	All comparison operators are supported (rather than just =)	— Subclause 7.7, “<joined table>”: Subfeature F041-1 extended so that the <join condition> is not limited to a <comparison predicate> with a <comp op> of <equals operator>
124	F051	Basic date and time	— Subclause 6.1, “<data type>”: <datetime type> including datetime literals, datetime comparisons, and datetime conversions
125	F051-01	DATE data type (including support of DATE literal)	— Subclause 5.3, “<literal>”: The <date literal> form of <datetime literal> — Subclause 6.1, “<data type>”: The DATE <datetime type> — Subclause 6.30, “<datetime value expression>”: For values of type DATE
126	F051-02	TIME data type (including support of TIME literal) with fractional seconds precision of at least 0.	— Subclause 5.3, “<literal>”: The <time literal> form of <datetime literal>, where the value of <unquoted time string> is simply <time value> that does not include the optional <time zone interval> — Subclause 6.1, “<data type>”: The TIME <datetime type> without the <with or without time zone> clause — Subclause 6.30, “<datetime value expression>”: For values of type TIME

	Feature ID	Feature Name	Feature Description
127	F051-03	TIMESTAMP data type (including support of TIMESTAMP literal) with fractional seconds precision of at least 0 and 6.	— Subclause 5.3, “<literal>”: The <timestamp literal> form of <datetime literal>, where the value of <unquoted timestamp string> is simply <time value> that does not include the optional <time zone interval> — Subclause 6.1, “<data type>”: The TIMESTAMP <datetime type> without the <with or without time zone> clause — Subclause 6.30, “<datetime value expression>”: For values of type TIMESTAMP
128	F051-04	Comparison predicate on DATE, TIME, and TIMESTAMP data types	— Subclause 8.2, “<comparison predicate>”: For comparison between values of the following types: DATE and DATE, TIME and TIME, TIMESTAMP and TIMESTAMP
129	F051-05	Explicit CAST between date-time types and character string types	— Subclause 6.12, “<cast specification>”: If support for Feature F201, “CAST function” is available, then CASTing between the following types: from character string to DATE, TIME, and TIMESTAMP; from DATE to DATE, TIMESTAMP, and character string; from TIME to TIME, TIMESTAMP, and character string; from TIMESTAMP to DATE, TIME, TIMESTAMP, and character string
130	F051-06	CURRENT_DATE	— Subclause 6.31, “<datetime value function>”: The <current date value function> — Subclause 6.30, “<datetime value expression>”: When the value is a <current date value function>
131	F051-07	LOCALTIME	— Subclause 6.31, “<datetime value function>”: The <current local time value function> — Subclause 6.30, “<datetime value expression>”: When the value is a <current local time value function> — Subclause 11.5, “<default clause>”: LOCALTIME option of <datetime value function>
132	F051-08	LOCALTIMESTAMP	— Subclause 6.31, “<datetime value function>”: The <current local timestamp value function> — Subclause 6.30, “<datetime value expression>”: When the value is a <current local timestamp value function> — Subclause 11.5, “<default clause>”: LOCALTIMESTAMP option of <datetime value function>
133	F081	UNION and EXCEPT in views	— Subclause 11.23, “<view definition>”: A <query expression> in a <view definition> may specify UNION, UNION ALL, and/or EXCEPT



	Feature ID	Feature Name	Feature Description
134	F131	<b>Grouped operations</b>	<p>— A <i>grouped view</i> is a view whose &lt;query expression&gt; is a &lt;query specification&gt; that contains an explicit or implicit &lt;group by clause&gt; or a &lt;having clause&gt; that is not contained in a &lt;subquery&gt;.</p> <p>Using this definition, Feature F131, “Grouped operations” is the union of its five subfeatures F131-01, F131-02, F131-03, F131-04, and F131-05.</p>
135	F131-01	WHERE, GROUP BY, and HAVING clauses supported in queries with grouped views	— Subclause 7.4, “<table expression>”: Even though a table in the <from clause> is a grouped view, the <where clause>, <group by clause>, and <having clause> may be specified
136	F131-02	Multiple tables supported in queries with grouped views	— Subclause 7.5, “<from clause>”: Even though a table in the <from clause> is a grouped view, the <from clause> may specify more than one <table reference>
137	F131-03	Set functions supported in queries with grouped views	— Subclause 7.12, “<query specification>”: Even though a table in the <from clause> is a grouped view, the <select list> may specify a <set function specification>
138	F131-04	Subqueries with GROUP BY and HAVING clauses and grouped views	— Subclause 7.15, “<subquery>”: A <subquery> in a <comparison predicate> is allowed to contain a <group by clause> and/or a <having clause> and/or it may identify a grouped view
139	F131-05	Single row SELECT with GROUP BY and HAVING clauses and grouped views	— Subclause 14.7, “<select statement: single row>”: The table in a <from clause> can be a grouped view — Subclause 14.7, “<select statement: single row>”: The <table expression> may specify a <group by clause> and/or a <having clause>



	Feature ID	Feature Name	Feature Description
140	F181	<b>Multiple module support</b>  NOTE 510 — The ability to associate multiple host compilation units with a single SQL-session at one time.	<p>— Subclause 13.1, “&lt;SQL-client module definition&gt;”: An SQL-agent can be associated with more than one &lt;SQL-client module definition&gt;. With this feature, it is possible to compile &lt;SQL-client module definition&gt;s or &lt;embedded SQL host program&gt;s separately and rely on the SQL-implementation to “link” the together properly at execution time. To ensure portability, applications should adhere to the following limitations:</p> <p>— Avoid linking modules having cursors with the same &lt;cursor name&gt;.</p> <p>— Avoid linking modules that prepare statements using the same &lt;SQL statement name&gt;.</p> <p>— Avoid linking modules that allocate descriptors with the same &lt;descriptor name&gt;.</p> <p>— Assume that the scope of an &lt;embedded exception declaration&gt; is a single compilation unit.</p> <p>— Assume that an &lt;embedded variable name&gt; can be referenced only in the same compilation unit in which it is declared.</p>
141	F201	<b>CAST function</b>  NOTE 511 — This means the support of CAST, where relevant, among all supported data types.	<p>— Subclause 6.12, “&lt;cast specification&gt;”: For all supported data types — Subclause 6.25, “&lt;value expression&gt;”: &lt;cast specification&gt;</p>
142	F221	<b>Explicit defaults</b>	<p>— Subclause 6.5, “&lt;contextually typed value specification&gt;”: &lt;default specification&gt;</p> <p>NOTE 512 — Including its use in UPDATE and INSERT statements.</p>
143	F261	<b>CASE expression</b>	<p>— Subclause 6.25, “&lt;value expression&gt;”: &lt;case expression&gt;</p>
144	F261-01	Simple CASE	<p>— Subclause 6.11, “&lt;case expression&gt;”: The &lt;simple case&gt; variation</p>
145	F261-02	Searched CASE	<p>— Subclause 6.11, “&lt;case expression&gt;”: The &lt;searched case&gt; variation</p>
146	F261-03	NULLIF	<p>— Subclause 6.11, “&lt;case expression&gt;”: The NULLIF &lt;case abbreviation&gt;</p>
147	F261-04	COALESCE	<p>— Subclause 6.11, “&lt;case expression&gt;”: The COALESCE &lt;case abbreviation&gt;</p>
148	F311	<b>Schema definition statement</b>	<p>— Subclause 11.1, “&lt;schema definition&gt;”</p>

	Feature ID	Feature Name	Feature Description
149	F311-01	CREATE SCHEMA	— Subclause 11.1, “<schema definition>”: Support for circular references in that <referential constraint definition>s in two different <table definition>s may reference columns in the other table
150	F311-02	CREATE TABLE for persistent base tables	— Subclause 11.1, “<schema definition>”: A <schema element> that is a <table definition> — Subclause 11.3, “<table definition>”: In the context of a <schema definition>
151	F311-03	CREATE VIEW	— Subclause 11.1, “<schema definition>”: A <schema element> that is a <view definition> — Subclause 11.23, “<view definition>”: In the context of a <schema definition> without the WITH CHECK OPTION clause and without support of Feature F081, “UNION and EXCEPT in views”
152	F311-04	CREATE VIEW: WITH CHECK OPTION	— Subclause 11.23, “<view definition>”: The WITH CHECK OPTION clause, in the context of a <schema definition>, but without support of Feature F081, “UNION and EXCEPT in views”
153	F311-05	GRANT statement	— Subclause 11.1, “<schema definition>”: A <schema element> that is a <grant statement> — Subclause 12.1, “<grant statement>”: In the context of a <schema definition>
154	F471	Scalar subquery values	— Subclause 6.25, “<value expression>”: A <value expression primary> can be a <scalar subquery>
155	F481	Expanded NULL predicate	— Subclause 8.8, “<null predicate>”: The <row value expression> can be something other than a <column reference>
156	F812	Basic flagging	— Part 1, Subclause 8.5, “SQL flagger”: With “level of flagging” specified to be Core SQL Flagging and “extent of checking” specified to be Syntax Only  NOTE 513 — This form of flagging identifies vendor extensions and other non-standard SQL by checking syntax only without requiring access to the catalog information.

	Feature ID	Feature Name	Feature Description
157	S011	<b>Distinct data types</b>	— Subclause 11.42, “<user-defined type definition>”: When <representation> is <predefined type> — Subclause 11.50, “<drop data type statement>”
158	T321	<b>Basic SQL-invoked routines</b>	— Subclause 11.51, “<SQL-invoked routine>” — Subclause 11.53, “<drop routine statement>” — If Feature T041, “Basic LOB data type support”, is supported, then the <locator indication> clause shall also be supported  NOTE 514 — “Routine” is the collective term for functions, methods, and procedures. This feature requires a conforming SQL-implementation to support both user-defined functions and user-defined procedures. An SQL-implementation that conforms to Core SQL shall support at least one language for writing routines; that language may be SQL. If the language is SQL, then the basic specification capability in Core SQL is the ability to specify a one-statement routine. Support for overloaded functions and procedures is not part of Core SQL.
159	T321-01	User-defined functions with no overloading	— Subclause 11.51, “<SQL-invoked routine>”: With <function specification>
160	T321-02	User-defined stored procedures with no overloading	— Subclause 11.51, “<SQL-invoked routine>”: With <SQL-invoked procedure>
161	T321-03	Function invocation	— Subclause 6.4, “<value specification> and <target specification>”: With a <value expression primary> that is a <routine invocation> — Subclause 10.4, “<routine invocation>”: For user-defined functions
162	T321-04	CALL statement	— Subclause 10.4, “<routine invocation>”: Used by <call statement>s — Subclause 16.1, “<call statement>”
163	T321-05	RETURN statement	— Subclause 16.2, “<return statement>”, if the SQL-implementation supports SQL routines
164	T631	<b>IN predicate with one list element</b>	— Subclause 8.4, “<in predicate>”: <in value list> containing exactly one <row value expression>
<sup>1</sup> A conforming SQL-implementation is required (by Clause 8, “Conformance”, in [ISO9075-1]) to support at least one embedded language or to support the SQL-client module binding for at least one host language.			

Table 37, “Feature taxonomy for optional features”, does not provide definitions of the features; the definition of those features is found in the Conformance Rules that are further summarized in Annex A, “SQL Conformance Summary”.

Table 37 — Feature taxonomy for optional features

	Feature ID	Feature Name
1	B011	Embedded Ada
2	B012	Embedded C
3	B013	Embedded COBOL
4	B014	Embedded Fortran
5	B015	Embedded MUMPS
6	B016	Embedded Pascal
7	B017	Embedded PL/I
8	B021	Direct SQL
9	B031	Basic dynamic SQL
10	B032	Extended dynamic SQL
11	B032-01	<describe input statement>
12	B033	Untyped SQL-invoked function arguments
13	B034	Dynamic specification of cursor attributes
14	B035	Non-extended descriptor names
15	B041	Extensions to embedded SQL exception declarations
16	B051	Enhanced execution rights
17	B111	Module language Ada
18	B112	Module language C
19	B113	Module language COBOL
20	B114	Module language Fortran
21	B115	Module language MUMPS
22	B116	Module language Pascal
23	B117	Module language PL/I
24	B121	Routine language Ada

	Feature ID	Feature Name
25	B122	Routine language C
26	B123	Routine language COBOL
27	B124	Routine language Fortran
28	B125	Routine language MUMPS
29	B126	Routine language Pascal
30	B127	Routine language PL/I
31	B128	Routine language SQL
32	F032	CASCADE drop behavior
33	F033	ALTER TABLE statement: DROP COLUMN clause
34	F034	Extended REVOKE statement
35	F034-01	REVOKE statement performed by other than the owner of a schema object
36	F034-02	REVOKE statement: GRANT OPTION FOR clause
37	F034-03	REVOKE statement to revoke a privilege that the grantee has WITH GRANT OPTION
38	F052	Intervals and datetime arithmetic
39	F053	OVERLAPS predicate
40	F111	Isolation levels other than SERIALIZABLE
41	F111-01	READ UNCOMMITTED isolation level
42	F111-02	READ COMMITTED isolation level
43	F111-03	REPEATABLE READ isolation level
44	F121	Basic diagnostics management
45	F121-01	GET DIAGNOSTICS statement
46	F121-02	SET TRANSACTION statement: DIAGNOSTICS SIZE clause
47	F122	Enhanced diagnostics management
48	F123	All diagnostics
49	F171	Multiple schemas per user

	Feature ID	Feature Name
50	F191	Referential delete actions
51	F200	TRUNCATE TABLE statement
52	F222	INSERT statement: DEFAULT VALUES clause
53	F231	Privilege tables
54	F231-01	TABLE_PRIVILEGES view
55	F231-02	COLUMN_PRIVILEGES view
56	F231-03	USAGE_PRIVILEGES view
57	F251	Domain support
58	F262	Extended CASE expression
59	F263	Comma-separated predicates in simple CASE expression
60	F271	Compound character literals
61	F281	LIKE enhancements
62	F291	UNIQUE predicate
63	F301	CORRESPONDING in query expressions
64	F302	INTERSECT table operator
65	F302-01	INTERSECT DISTINCT table operator
66	F302-02	INTERSECT ALL table operator
67	F304	EXCEPT ALL table operator
68	F312	MERGE statement
69	F313	Enhanced MERGE statement
70	F321	User authorization
71	F341	Usage tables
72	F361	Subprogram support
73	F381	Extended schema manipulation
74	F381-01	ALTER TABLE statement: ALTER COLUMN clause

	<b>Feature ID</b>	<b>Feature Name</b>
<b>75</b>	F381-02	ALTER TABLE statement: ADD CONSTRAINT clause
<b>76</b>	F381-03	ALTER TABLE statement: DROP CONSTRAINT clause
<b>77</b>	<b>F382</b>	<b>Alter column data type</b>
<b>78</b>	<b>F391</b>	<b>Long identifiers</b>
<b>79</b>	<b>F392</b>	<b>Unicode escapes in identifiers</b>
<b>80</b>	<b>F393</b>	<b>Unicode escapes in literals</b>
<b>81</b>	<b>F394</b>	<b>Optional normal form specification</b>
<b>82</b>	<b>F401</b>	<b>Extended joined table</b>
<b>83</b>	F401-01	NATURAL JOIN
<b>84</b>	F401-02	FULL OUTER JOIN
<b>85</b>	F401-04	CROSS JOIN
<b>86</b>	<b>F402</b>	<b>Named column joins for LOBs, arrays, and multisets</b>
<b>87</b>	<b>F403</b>	<b>Partitioned join tables</b>
<b>88</b>	<b>F411</b>	<b>Time zone specification</b>
<b>89</b>	<b>F421</b>	<b>National character</b>
<b>90</b>	<b>F431</b>	<b>Read-only scrollable cursors</b>
<b>91</b>	F431-01	FETCH with explicit NEXT
<b>92</b>	F431-02	FETCH FIRST
<b>93</b>	F431-03	FETCH LAST
<b>94</b>	F431-04	FETCH PRIOR
<b>95</b>	F431-05	FETCH ABSOLUTE
<b>96</b>	F431-06	FETCH RELATIVE
<b>97</b>	<b>F441</b>	<b>Extended set function support</b>
<b>98</b>	<b>F442</b>	<b>Mixed column references in set functions</b>
<b>99</b>	<b>F451</b>	<b>Character set definition</b>

	<b>Feature ID</b>	<b>Feature Name</b>
<b>100</b>	<b>F461</b>	<b>Named character sets</b>
<b>101</b>	<b>F491</b>	<b>Constraint management</b>
<b>102</b>	<b>F502</b>	<b>Enhanced documentation tables</b>
<b>103</b>	F502-01	SQL_SIZING_PROFILES view
<b>104</b>	F502-02	SQL_IMPLEMENTATION_INFO view
<b>105</b>	F502-03	SQL_PACKAGES view
<b>106</b>	<b>F521</b>	<b>Assertions</b>
<b>107</b>	<b>F531</b>	<b>Temporary tables</b>
<b>108</b>	<b>F555</b>	<b>Enhanced seconds precision</b>
<b>109</b>	<b>F561</b>	<b>Full value expressions</b>
<b>110</b>	<b>F571</b>	<b>Truth value tests</b>
<b>111</b>	<b>F591</b>	<b>Derived tables</b>
<b>112</b>	<b>F611</b>	<b>Indicator data types</b>
<b>113</b>	<b>F641</b>	<b>Row and table constructors</b>
<b>114</b>	<b>F651</b>	<b>Catalog name qualifiers</b>
<b>115</b>	<b>F661</b>	<b>Simple tables</b>
<b>116</b>	<b>F671</b>	<b>Subqueries in CHECK</b>
<b>117</b>	<b>F672</b>	<b>Retrospective check constraints</b>
<b>118</b>	<b>F690</b>	<b>Collation support</b>
<b>119</b>	<b>F692</b>	<b>Enhanced collation support</b>
<b>120</b>	<b>F693</b>	<b>SQL-session and client module collations</b>
<b>121</b>	<b>F695</b>	<b>Translation support</b>
<b>122</b>	<b>F701</b>	<b>Referential update actions</b>
<b>123</b>	<b>F711</b>	<b>ALTER domain</b>
<b>124</b>	<b>F721</b>	<b>Deferrable constraints</b>



	Feature ID	Feature Name
125	F731	INSERT column privileges
126	F741	Referential MATCH types
127	F751	View CHECK enhancements
128	F761	Session management
129	F762	CURRENT_CATALOG
130	F763	CURRENT_SCHEMA
131	F771	Connection management
132	F781	Self-referencing operations
133	F791	Insensitive cursors
134	F801	Full set function
135	F813	Extended flagging — Part 1, Subclause 8.5, “SQL flagger”: With “level of flagging” specified to be Core SQL Flagging and “extent of checking” specified to be Catalog Lookup
136	F821	Local table references
137	F831	Full cursor update
138	F831-01	Updateable scrollable cursors
139	F831-02	Updateable ordered cursors
140	F841	LIKE_REGEX predicate
141	F842	OCCURRENCES_REGEX function
142	F843	POSITION_REGEX function
143	F844	SUBSTRING_REGEX function
144	F845	TRANSLATE_REGEX function
145	F846	Octet support in regular expression operators
146	F847	Nonconstant regular expressions
147	F850	Top-level <order by clause> in <query expression>
148	F851	<order by clause> in subqueries

	Feature ID	Feature Name
149	F852	Top-level <order by clause> in views
150	F855	Nested <order by clause> in <query expression>
151	F856	Nested <fetch first clause> in <query expression>
152	F857	Top-level <fetch first clause> in <query expression>
153	F858	<fetch first clause> in subqueries
154	F859	Top-level <fetch first clause> in views
155	S023	Basic structured types
156	S024	Enhanced structured types
157	S025	Final structured types
158	S026	Self-referencing structured types
159	S027	Create method by specific method name
160	S028	Permutable UDT options list
161	S041	Basic reference types
162	S043	Enhanced reference types
163	S051	Create table of type
164	S071	SQL paths in function and type name resolution
165	S081	Subtables
166	S091	Basic array support
167	S091-01	Arrays of built-in data types
168	S091-02	Arrays of distinct types
169	S091-03	Array expressions
170	S092	Arrays of user-defined types
171	S094	Arrays of reference types
172	S095	Array constructors by query
173	S096	Optional array bounds

	Feature ID	Feature Name
174	S097	Array element assignment
<del>175</del>	S098	ARRAY_AGG
176	S111	ONLY in query expressions
177	S151	Type predicate
178	S161	Subtype treatment
179	S162	Subtype treatment for references
180	S201	SQL-invoked routines on arrays
181	S201-01	Array parameters
182	S201-02	Array as result type of functions
183	S202	SQL-invoked routines on multisets
184	S211	User-defined cast functions
185	S231	Structured type locators
186	S232	Array locators
187	S233	Multiset locators
188	S241	Transform functions
189	S242	Alter transform statement
190	S251	User-defined orderings
191	S261	Specific type method
192	S271	Basic multiset support
193	S272	Multisets of user-defined types
194	S274	Multisets of reference types
195	S275	Advanced multiset support
196	S281	Nested collection types
197	S291	Unique constraint on entire row
198	S301	Enhanced UNNEST

	Feature ID	Feature Name
199	T011	Timestamp in Information Schema
200	T021	BINARY and VARBINARY data types
201	T022	Advanced support for BINARY and VARBINARY data types
202	T023	Compound binary literals
203	T024	Spaces in binary literals
204	T031	BOOLEAN data type
205	T041	Basic LOB data type support
206	T041-01	<p>BLOB data type</p> <ul style="list-style-type: none"> <li>— Subclause 5.2, “&lt;token&gt; and &lt;separator&gt;”: The &lt;reserved word&gt;s BINARY, BLOB, LARGE, and OBJECT</li> <li>— Subclause 5.3, “&lt;literal&gt;”: &lt;binary string literal&gt;</li> <li>— Subclause 6.1, “&lt;data type&gt;”: The BINARY LARGE OBJECT data type</li> <li>— Subclause 6.28, “&lt;string value expression&gt;”: For values of type BINARY LARGE OBJECT</li> <li>— Subclause 13.6, “Data type correspondences”: Type correspondences for BINARY LARGE OBJECT for all supported languages</li> </ul>
207	T041-02	<p>CLOB data type</p> <ul style="list-style-type: none"> <li>— Subclause 5.2, “&lt;token&gt; and &lt;separator&gt;”: The &lt;reserved word&gt;s CHARACTER, CLOB, LARGE, and OBJECT</li> <li>— Subclause 6.1, “&lt;data type&gt;”: The CHARACTER LARGE OBJECT data type</li> <li>— Subclause 6.28, “&lt;string value expression&gt;”: For values of type CHARACTER LARGE OBJECT</li> <li>— Subclause 13.6, “Data type correspondences”: Type correspondences for CHARACTER LARGE OBJECT for all supported languages</li> <li>— The implicit casting among the fixed-length and variable-length character string types supported by subfeature E021-10 is extended to support the character large object type</li> </ul>

	Feature ID	Feature Name
208	T041-03	<p>POSITION, LENGTH, LOWER, TRIM, UPPER, and SUBSTRING functions for LOB data types</p> <p>— Subclause 6.27, “&lt;numeric value function&gt;”: The &lt;position expression&gt; for expressions of type BINARY LARGE OBJECT and CHARACTER LARGE OBJECT</p> <p>— Subclause 6.27, “&lt;numeric value function&gt;”: The &lt;char length expression&gt; for expressions of type CHARACTER LARGE OBJECT</p> <p>— Subclause 6.27, “&lt;numeric value function&gt;”: The &lt;octet length expression&gt; for expressions of type BINARY LARGE OBJECT and CHARACTER LARGE OBJECT</p> <p>— Subclause 6.29, “&lt;string value function&gt;”: The &lt;fold&gt; function for expressions of type CHARACTER LARGE OBJECT</p> <p>— Subclause 6.29, “&lt;string value function&gt;”: The &lt;trim function&gt; for expressions of type CHARACTER LARGE OBJECT</p> <p>— Subclause 6.29, “&lt;string value function&gt;”: The &lt;binary trim function&gt; for expressions of type BINARY LARGE OBJECT</p> <p>— Subclause 6.29, “&lt;string value function&gt;”: The &lt;character substring function&gt; for expressions of type CHARACTER LARGE OBJECT</p> <p>— Subclause 6.29, “&lt;string value function&gt;”: The &lt;binary substring function&gt; for expressions of type BINARY LARGE OBJECT</p>
209	T041-04	<p>Concatenation of LOB data types</p> <p>— Subclause 6.28, “&lt;string value expression&gt;”: The &lt;concatenation&gt; expression for expressions of type CHARACTER LARGE OBJECT</p> <p>— Subclause 6.28, “&lt;string value expression&gt;”: The &lt;binary concatenation&gt; expression for expressions of type BINARY LARGE OBJECT</p>
210	T041-05	<p>LOB locator: non-holdable</p> <p>— Subclause 13.3, “&lt;externally-invoked procedure&gt;”: &lt;locator indication&gt;</p> <p>— Subclause 14.17, “&lt;free locator statement&gt;”</p>
211	T042	<b>Extended LOB data type support</b>
212	T043	<b>Multiplier T</b>
213	T044	<b>Multiplier P</b>
214	T051	<b>Row types</b>
215	T052	<b>MAX and MIN for row types</b>
216	T053	<b>Explicit aliases for all-fields reference</b>
217	T061	<b>UCS support</b>
218	T071	<b>BIGINT data type</b>
219	T101	<b>Enhanced nullability determination</b>

	Feature ID	Feature Name
220	T111	Updatable joins, unions, and columns
221	T121	WITH (excluding RECURSIVE) in query expression
222	T122	WITH (excluding RECURSIVE) in subquery
223	T131	Recursive query
224	T132	Recursive query in subquery
225	T141	SIMILAR predicate
226	T151	DISTINCT predicate
227	T152	DISTINCT predicate with negation
228	T171	LIKE clause in table definition
229	T172	AS subquery clause in table definition
230	T173	Extended LIKE clause in table definition
231	T174	Identity columns
232	T175	Generated columns
233	T176	Sequence generator support
234	T191	Referential action RESTRICT
235	T201	Comparable data types for referential constraints
236	T211	Basic trigger capability
237	T211-01	Triggers activated on UPDATE, INSERT, or DELETE of one base table.
238	T211-02	BEFORE triggers
239	T211-03	AFTER triggers
240	T211-04	FOR EACH ROW triggers
241	T211-05	Ability to specify a search condition that shall be <i>True</i> before the trigger is invoked.
242	T211-06	Support for run-time rules for the interaction of triggers and constraints.
243	T211-07	TRIGGER privilege
244	T211-08	Multiple triggers for the same event are executed in the order in which they were created in the catalog.

	Feature ID	Feature Name
245	T212	Enhanced trigger capability
246	T213	INSTEAD OF triggers
247	T231	Sensitive cursors
248	T241	START TRANSACTION statement
249	T251	SET TRANSACTION statement: LOCAL option
250	T261	Chained transactions
251	T271	Savepoints
252	T272	Enhanced savepoint management
253	T281	SELECT privilege with column granularity
254	T285	Enhanced derived column names
255	T301	Functional dependencies
256	T312	OVERLAY function
257	T322	Overloading of SQL-invoked functions and procedures
258	T323	Explicit security for external routines
259	T324	Explicit security for SQL routines
260	T325	Qualified SQL parameter references
261	T326	Table functions
262	T331	Basic roles
263	T332	Extended roles
264	T351	Bracketed comments
265	T431	Extended grouping capabilities
266	T432	Nested and concatenated GROUPING SETS
267	T433	Multiargument GROUPING function
268	T434	GROUP BY DISINCT
269	T441	ABS and MOD functions

	Feature ID	Feature Name
270	T461	Symmetric BETWEEN predicate
271	T471	Result sets return value
272	T491	LATERAL derived table
273	T501	Enhanced EXISTS predicate
274	T511	Transaction counts
275	T551	Optional key words for default syntax
276	T561	Holdable locators
277	T571	Array-returning external SQL-invoked functions
278	T572	Multiset-returning external SQL-invoked functions
279	T581	Regular expression substring function
280	T591	UNIQUE constraints of possibly null columns
281	T601	Local cursor references
282	T611	Elementary OLAP operations
283	T612	Advanced OLAP operations
284	T613	Sampling
285	T621	Enhanced numeric functions
286	T641	Multiple column assignment
287	T651	SQL-schema statements in SQL routines
288	T652	SQL-dynamic statements in SQL routines
289	T653	SQL-schema statements in external routines
290	T654	SQL-dynamic statements in external routines
291	T655	Cyclically dependent routines



## Annex G

### (informative)

### Defect Reports not addressed in this edition of ISO/IEC 9075

Each entry in this Annex describes a reported defect in the previous edition of this part of ISO/IEC 9075 that remains in this edition.

1) **Subclause 10.4, “<routine invocation>”:**

There is no definition of how to pass values of type BOOLEAN or of large object types as arguments to invocations of external routines. More generally, the question of how to convert a value of any SQL type to a value of an appropriate host language type at the interface to an SQL-invoked routine is not addressed. The rules in Subclause 13.4, “Calls to an <externally-invoked procedure>”, are appropriate, but they are not referenced by the rules of Subclause 10.4, “<routine invocation>”.

2) **Subclause 21.1, “<embedded SQL host program>”:**

SR 21)h)i)6) and SR 21)l)i)3)B)VI) both refer to the SQL data type that corresponds to a given host language data type, as determined by application of the rules in Subclause 13.6, “Data type correspondences”. These two syntax rules are sometimes ambiguous, because Subclause 13.6, “Data type correspondences” does not always give exactly one SQL data type for a given host language type, as can be seen by inspection of the data type correspondence tables given in that Subclause. For example, Table 17, “Data type correspondences for C”, in which the C data type “pointer to long” maps to both INTEGER and BOOLEAN.

*(Blank page)*

## Bibliography

- [1] [ISO9075-3] ISO/IEC 9075-3:2008, *Information technology — Database languages — SQL — Part 3: Call-Level Interface (SQL/CLI)*.
- [2] [ISO9075-4] ISO/IEC 9075-4:2008, *Information technology — Database languages — SQL — Part 4: Persistent Stores Modules (SQL/PSM)*.
- [3] [ISO9075-9] ISO/IEC 9075-9:2008, *Information technology — Database languages — SQL — Part 9: Management of External Data (SQL/MED)*.
- [4] [ISO9075-10] ISO/IEC 9075-10:2008, *Information technology — Database languages — SQL — Part 10: Object Language Bindings (SQL/OLB)*.
- [5] [ISO9075-13] ISO/IEC 9075-13:2008, *Information technology — Database languages — SQL — Part 13: SQL Routines and Types Using the Java™ Programming Language (SQL/JRT)*.
- [6] [ISO9075-14] ISO/IEC 9075-14:2008, *Information technology — Database languages — SQL — Part 14: XML-Related Specifications (SQL/XML)*.
- [7] [POSIX] ISO/IEC 9945-1:2002, *Standard for Information Technology — Portable Operating System Interface (POSIX) — Part 1: Base definitions*.

*(Blank page)*

## Index

Index entries appearing in **boldface** indicate the page where the word, phrase, or BNF nonterminal was defined; index entries appearing in *italics* indicate a page where the BNF nonterminal was used in a Format; and index entries appearing in roman type indicate a page where the word, phrase, or BNF nonterminal was used in a heading, Function, Syntax Rule, Access Rule, General Rule, Leveling Rule, Table, or other descriptive text.

### — A —

A • 146, 448

abandoned • 790, 791, 792, 793, 796, 797, 799, 800, 801, 802

ABS • 147, 257, 266, 298, 1218

ABSOLUTE • 146, 857, 905, 906

<absolute value expression> • 32, 256, **257**, 260, 262, 266, 1218

abstract character • 5

access mode • 124, 127, 131, 524, 537, 836, 869, 871, 876, 881, 889, 908, 910, 945, 948, 956, 958, 962, 963, 970, 1121, 1149

ACTION • 69, 70, 146, 597, 599, 836, 934, 935, 936, 937

<action> • 119, 120, 121, 561, 772, 774, 776, 777, **779**, 781, 786, 787, 802, 803, 1166, 1181, 1187, 1189, 1193, 1212, 1214, 1262, 1263

activated by • 136, 913, 920, 926, 932, 939

active • 95, 104, 113, 123, 128, 130, 131, 132, 135, 137, 398, 518, 519, 525, 835, 836, 903, 945, 947, 951, 955, 957, 958, 966, 971, 972, 1037, 1121, 1135, 1136, 1149, 1228, 1242, 1250

active SQL-connection • 128

*active SQL-transaction* • 945, 947, 966, 971, 972, 1121, 1149

<actual identifier> • **161**

ADA • 146, 489, 513, 529, 731, 738, 741, 808, 810, 812, 814, 825, 828, 829, 1059, 1163, 1164

<Ada array locator variable> • **1072**, 1075, 1077, 1197

<Ada assignment operator> • **1071**

<Ada BINARY variable> • **1072**, 1073, 1077, 1202

<Ada BLOB locator variable> • **1072**, 1074, 1076, 1205

<Ada BLOB variable> • **1072**, 1074, 1076, 1205

<Ada CLOB locator variable> • **1072**, 1075, 1076, 1205

<Ada CLOB variable> • **1072**, 1074, 1076, 1205

<Ada derived type specification> • *1071*, **1072**, 1073

<Ada host identifier> • *1056*, **1071**, 1072, 1073, 1074, 1075

<Ada initial value> • **1071**, 1076

<Ada multiset locator variable> • **1072**, 1075, 1077, 1198

<Ada qualified type specification> • **1071**, 1073, 1077, 1209

<Ada REF variable> • **1072**, 1075, 1077, 1190

<Ada type specification> • **1071**

<Ada unqualified type specification> • **1071**, 1077, 1209

<Ada user-defined type locator variable> • **1072**, 1075, 1077, 1196

<Ada user-defined type variable> • **1072**, 1074, 1077, 1199

<Ada VARBINARY variable> • **1072**, 1073, 1077, 1203

<Ada variable definition> • *1056*, **1071**, 1073, 1075, 1076

ADD • 146, 604, 609, 618, 644, 647, 693, 697, 703, 759

<add attribute definition> • 88, 692, **693**, 694

<add column definition> • 603, **604**, 1267

<add column scope clause> • 606, **609**, 1171, 1191

<add domain constraint definition> • 641, **644**, 1181

<add original method specification> • 692, **697**, 702

<add overriding method specification> • 692, **703**, 707

<add table constraint definition> • 596, 603, **618**, 1171

<add transform element list> • 758, **759**, 760

additional parameter • 88

*additional result sets returned* • 1150

ADMIN • 122, 146, 781, 782, 783, 786, 790, 791, 800, 1217

AFTER • 74, 133, 134, 136, 146, 257, 520, 666, 939

AFTER trigger • 133

<aggregate function> • 205, 206, 207, 209, 210, 230, 252, 372, 481, 485, **547**, 548, 557, 558, 1174, 1175, 1249

aggregated argument • 79, 203, 205, 355, 358, 365, 371, 374, 619, 1214

*aggregated column reference* • 205

aggregation query • 203, 205, 371, 374, 549

ALL • 50, 64, 65, 147, 274, 277, 285, 308, 309, 341, 342, 346, 376, 379, 380, 381, 382, 384, 385, 388, 390, 391, 430, 547, 548, 555, 623, 633, 713, 763, 764, 777, 779, 780, 802, 915, 924, 927, 931, 951, 955, 966, 1126, 1169, 1243, 1250, 1251

- <all> • **430**
- <all fields column name list> • **366**, 368, 374, 1208
- <all fields reference> • **366**, 367, 368, 374, 1208
- <all info target> • **1126**, 1140
- <all information> • 1125, **1126**, 1139, 1141, 1167
- <all qualifier> • **1126**, 1140
- ALLOCATE • 147, 993, 1036
- <allocate cursor statement> • 10, 86, 93, 99, 101, 102, 107, 111, 113, 115, 117, 170, 538, 833, 853, 854, **1036**, 1037, 1038, 1130, 1162, 1218, 1242
- <allocate descriptor statement> • 86, 109, 111, 170, 833, **993**, 994, 998, 1130, 1160, 1241
- ALTER • 147, 603, 606, 611, 617, 620, 624, 637, 641, 642, 643, 644, 645, 647, 664, 692, 740, 754, 758, 768, 801, 802, 1165, 1180, 1181
- <alter column action> • **606**
- <alter column data type clause> • 606, **612**, 614, 1171
- <alter column definition> • 589, 603, **606**, 607, 608, 609, 610, 612, 615, 1170, 1211
- <alter domain action> • **641**
- <alter domain statement> • 105, 167, 589, 601, 638, **641**, 642, 643, 644, 645, 801, 832, 1130, 1180
- <alter group> • **758**, 759, 761
- <alter identity column option> • **615**
- <alter identity column specification> • 606, **615**, 1211
- <alter routine behavior> • **740**
- <alter routine characteristic> • **740**
- <alter routine characteristics> • **740**, 741
- <alter routine statement> • 88, 106, **740**, 742, 831, 1130, 1171
- <alter sequence generator option> • **768**
- <alter sequence generator options> • 507, 615, **768**
- <alter sequence generator restart option> • 507, 508, 615, **768**
- <alter sequence generator statement> • 82, 106, **768**, 832, 1130, 1212
- <alter table action> • **603**
- <alter table statement> • 105, 167, 585, 586, 587, 593, 595, 596, 597, 601, **603**, 604, 605, 606, 609, 610, 611, 612, 616, 617, 618, 619, 620, 624, 637, 647, 664, 754, 801, 802, 831, 1130, 1267
- <alter transform action> • **758**
- <alter transform action list> • **758**
- <alter transform statement> • 106, **758**, 759, 761, 832, 1130, 1200
- <alter type action> • **692**
- <alter type statement> • 88, 106, **692**, 693, 694, 695, 697, 703, 708, 832, 1130, 1188
- ALWAYS • 60, 146, 576, 584, 587
- ambiguous • 133, 367, 1011, 1048, 1049, 1050, 1052, 1053, 1287
- ambiguous cursor name* • 1011, 1144
- <ampersand> • 139, **140**, 144, 149, 152, 156, 1055, 1057
- anchor expression • 378, 383
- anchor name • 378
- AND • 20, 33, 73, 75, 127, 147, 208, 270, 283, 299, 300, 302, 339, 357, 358, 411, 440, 556, 595, 751, 805, 807, 808, 878, 955, 956, 957, 958, 1056, 1163
- AND-component • 73, 77
- ANY • 65, 147, 413, 430, 547, 549, 552, 557, 1204
- apparent value • 31
- applicable for • 122, 783, 1019
- applicable from-sql function • 500, 501, 1030
- applicable to-sql function • 502, 503, 727, 1026
- appropriate user-defined cast function • 51
- <approximate numeric literal> • 31, **153**, 157, 159, 221, 223, 556, 1229, 1258
- <approximate numeric type> • 30, 31, **174**, 179, 183, 1021, 1230, 1231, 1258
- approximate numeric types • 15, 1232
- arbitrary character specifier • 416, 418
- arbitrary octet specifier • 418
- arbitrary string specifier • 416, 418
- ARE • 147, 810, 1056
- ARRAY • 13, 47, 48, 49, 99, 147, 175, 183, 196, 197, 219, 306, 327, 394, 395, 470, 492, 498, 589, 824, 840, 842, 844, 845, 846, 847, 875, 880, 894, 985, 986, 998, 1000, 1007, 1021, 1024, 1030, 1031, 1105, 1106, 1193
- array • iv, iv, v, 13, 17, 48, 49, 65, 96, 97, 129, 179, 183, 184, 192, 197, 219, 248, 251, 266, 304, 305, 306, 307, 327, 333, 372, 405, 442, 454, 457, 458, 462, 465, 489, 491, 492, 494, 498, 530, 532, 533, 534, 535, 536, 538, 551, 557, 589, 680, 681, 698, 699, 722, 723, 737, 738, 811, 812, 822, 826, 829, 859, 860, 864, 875, 880, 894, 895, 896, 899, 900, 985, 986, 1021, 1029, 1060, 1062, 1084, 1115, 1146, 1155, 1156, 1193, 1194, 1195, 1197, 1200, 1201, 1246, 1249, 1280
- <array aggregate function> • 252, 547, **548**, 551, 556, 557, 558, 1194, 1236, 1249
- <array concatenation> • 49, **304**, 305
- array data, right truncation* • 219, 305, 307, 457, 462, 551, 557, 1144, 1150
- <array element> • **306**, 307, 1006, 1007
- array element error* • 248, 490, 536, 860, 865, 896, 1145
- <array element list> • **306**, 307, 1006, 1007
- <array element reference> • 49, 188, 189, **248**, 372, 1005, 1193
- array locator parameter • 812
- array locator variable • 1075, 1083, 1091, 1098, 1103, 1108, 1115
- <array primary> • **304**
- <array type> • 48, **175**, 179, 184, 219, 688, 721, 813, 986, 1072, 1075, 1080, 1083, 1088, 1091, 1095, 1098, 1101,

- 1103, 1106, 1108, 1112, 1115, 1193, 1194, 1197, 1219, 1231
- <array value constructor> • 188, 189, **306**, 307, 1247
- <array value constructor by enumeration> • **306**, 307, 1193
- <array value constructor by query> • 251, **306**, 307, 391, 1185, 1194, 1247
- <array value expression> • 248, 250, 252, **304**, 305, 1193
- <array value expression 1> • **304**
- array-ordered • 17, 343, 481, 484, 485, 1172
- ARRAY\_AGG • 65, 147, 548, 1255
- AS • 96, 147, 196, 208, 209, 215, 217, 218, 219, 220, 221, 224, 225, 226, 227, 228, 233, 235, 236, 249, 291, 292, 294, 295, 296, 308, 309, 311, 312, 325, 326, 327, 328, 329, 337, 350, 351, 352, 356, 366, 368, 369, 370, 372, 376, 386, 389, 396, 409, 456, 461, 516, 533, 555, 556, 574, 576, 582, 584, 625, 626, 631, 638, 648, 666, 674, 675, 677, 679, 680, 681, 686, 688, 691, 697, 698, 699, 700, 713, 714, 715, 722, 744, 745, 746, 747, 766, 825, 828, 866, 868, 878, 883, 886, 888, 893, 944, 961, 969, 1025, 1029, 1060, 1061, 1062, 1064, 1072, 1074, 1075, 1079, 1080, 1082, 1083, 1084, 1088, 1090, 1091, 1095, 1097, 1098, 1101, 1102, 1103, 1105, 1106, 1107, 1108, 1112, 1114, 1115, 1126, 1134, 1135, 1139, 1188, 1211
- <as clause> • **366**, 370, 371, 375, 386, 1214, 1248, 1261
- <as subquery clause> • 55, 230, 573, **574**, 576, 577, 580, 582, 1211
- ASC • 61, 62, 146, 559
- ascending sequence generator • 82, 505
- ASENSITIVE • 94, 100, 102, 147, 851, 852, 1038, 1183, 1213
- asensitive • 102, 103, 1227, 1246
- ASSERTION • 146, 571, 611, 620, 624, 637, 661, 663, 664, 748, 754, 801
- <assertion definition> • 105, 203, 332, 567, **661**, 662, 831, 1130, 1176, 1179
- assignable • **7**
- <assigned row> • 794, 795, 796, 798, 799, **892**, 893, 896, 1183, 1189
- ASSIGNMENT • 146, 686, 745, 746
- assignment • **7**
- associated array type • 1075, 1083, 1091, 1098, 1103, 1108, 1115
- associated multiset type • 1075, 1083, 1091, 1098, 1103, 1108, 1115
- associated SQL data type • 1074, 1082, 1090, 1097, 1102, 1107, 1114
- associated user-defined type • 986, 1058, 1062, 1065, 1074, 1075, 1082, 1083, 1090, 1091, 1097, 1102, 1103, 1107, 1108, 1114
- associated value • 38, 39, 1024
- associated with • 19, 28, 38, 39, 41, 45, 48, 56, 60, 61, 68, 82, 84, 85, 86, 93, 95, 97, 123, 125, 128, 129, 130, 137, 151, 168, 194, 196, 246, 261, 333, 367, 368, 388, 389, 503, 524, 526, 531, 533, 539, 580, 617, 648, 669, 726, 729, 730, 731, 732, 733, 735, 761, 762, 763, 765, 792, 808, 812, 822, 834, 835, 897, 911, 955, 962, 965, 989, 990, 991, 992, 1015, 1017, 1021, 1029, 1035, 1039, 1120, 1121, 1134, 1143, 1171, 1225, 1228, 1230, 1235, 1243, 1245, 1257, 1271
- <asterisk> • 23, 79, 139, **140**, 149, 254, 255, 293, 346, 366, 368, 369, 371, 420, 422, 423, 432, 547, 1006, 1094, 1219
- <asterisked identifier> • **366**, 367
- <asterisked identifier chain> • 346, **366**, 367, 374, 1215
- ASYMMETRIC • 147, 411, 1218
- AT • 147, 288, 296
- ATOMIC • 147, 666, 1064
- atomic • **6**, 112, 113, 122, 132, 834, 837, 955, 957, 958
- atomic execution context • 113
- atomic SQL-statement • 112, 113, 837
- attempt to assign to non-updatable column • 1047, 1144
- attempt to assign to ordering column • 1046, 1144
- attempt to replace a zero-length string • 496, 1145
- attempt to return too many result sets • 536, 1150
- ATTRIBUTE • 146, 693, 695
- attribute • **7**, 52
- <attribute default> • 53, **690**, 691, 1188
- <attribute definition> • 589, 674, 677, 687, **690**, 691, 693, 694, 1179, 1187, 1189
- <attribute name> • **163**, 169, 171, 243, 578, 629, 630, 675, 677, 678, 690, 695, 1186
- <attribute or method reference> • 188, 189, **241**, 242, 1190
- ATTRIBUTES • 146, 1003
- <attributes specification> • **1003**, 1013, 1163
- <attributes variable> • 10, 854, 902, **1003**, 1004, 1013
- augmented SQL parameter declaration list • 681, 684, 701, 705, 706, 710, 719, 720, 724
- AUTHORIZATION • 147, 567, 568, 805, 971, 1056, 1237
- <authorization identifier> • 81, 118, 119, 120, 121, **161**, 168, 169, 203, 217, 218, 230, 244, 246, 280, 332, 517, 540, 544, 561, 563, 567, 568, 569, 570, 575, 581, 582, 585, 602, 603, 617, 623, 629, 630, 636, 639, 641, 646, 648, 649, 650, 652, 654, 657, 659, 661, 664, 668, 669, 671, 684, 692, 713, 732, 733, 735, 736, 741, 744, 745, 746, 748, 751, 753, 756, 758, 760, 764, 767, 768, 769, 779, 780, 782, 785, 791, 805, 812, 868, 871, 875, 880, 887, 888, 971, 1046, 1053, 1228, 1229, 1237
- authorization stack • 91, 118, 131, 525, 526, 538, 813, 941, 963, 1033, 1120, 1122, 1236
- AVG • 64, 147, 547, 549, 552, 1231, 1236

## — B —

- Feature B011, "Embedded Ada" • 1076, 1153, 1159
- Feature B012, "Embedded C" • 1085, 1153, 1159

- Feature B013, "Embedded COBOL" • 1092, 1153, 1159
- Feature B014, "Embedded Fortran" • 1099, 1153, 1159
- Feature B015, "Embedded MUMPS" • 1104, 1153, 1159
- Feature B016, "Embedded Pascal" • 1109, 1153, 1159, 1160
- Feature B017, "Embedded PL/I" • 1116, 1153, 1160
- Feature B021, "Direct SQL" • 1122, 1160
- Feature B031, "Basic dynamic SQL" • 171, 194, 994, 995, 998, 1002, 1013, 1022, 1026, 1031, 1033, 1034, 1035, 1040, 1041, 1042, 1043, 1045, 1047, 1160, 1161, 1162
- Feature B032, "Extended dynamic SQL" • 171, 994, 1016, 1022, 1033, 1038, 1049, 1053, 1162
- Feature B033, "Untyped SQL-invoked function arguments" • 538, 1162
- Feature B034, "Dynamic specification of cursor attributes" • 1013, 1162, 1163
- Feature B035, "Non-extended descriptor names" • 171, 1163
- Feature B041, "Extensions to embedded SQL exception declarations" • 1070, 1163
- Feature B051, "Enhanced execution rights" • 808, 1065, 1163
- Feature B111, "Module language Ada" • 808, 1153, 1163
- Feature B112, "Module language C" • 808, 1153, 1163
- Feature B113, "Module language COBOL" • 808, 1153, 1163
- Feature B114, "Module language Fortran" • 809, 1153, 1163
- Feature B115, "Module language MUMPS" • 809, 1153, 1164
- Feature B116, "Module language Pascal" • 809, 1153, 1164
- Feature B117, "Module language PL/I" • 809, 1153, 1164
- Feature B121, "Routine language Ada" • 738, 1153, 1164
- Feature B122, "Routine language C" • 739, 1153, 1164
- Feature B123, "Routine language COBOL" • 739, 1153, 1164
- Feature B124, "Routine language Fortran" • 739, 1153, 1164
- Feature B125, "Routine language MUMPS" • 739, 1153, 1164
- Feature B126, "Routine language Pascal" • 739, 1153, 1164, 1165
- Feature B127, "Routine language PL/I" • 739, 1153, 1165
- Feature B128, "Routine language SQL" • 739, 1154, 1165
- base column • 61
- based on • 16, 52, 70, 71, 90, 184, 252, 290, 408, 553, 617, 629, 690, 693, 695, 712, 713, 822, 1036, 1194, 1202
- <basic identifier chain> • **198**, 200, 201, 202, 204
- <basic sequence generator option> • 615, **766**, 768
- basis • 90, 198, 199, 200, 201, 202, 203, 367, 368, 629
- basis length • 198, 199, 200, 367
- basis table • 628
- BEFORE • 74, 133, 134, 136, 146, 200, 666, 668, 932, 941
- BEFORE trigger • 133
- BEGIN • 147, 666, 1056, 1064
- BERNOULLI • 146, 325, 333
- BETWEEN • 147, 208, 357, 358, 411, 1218
- <between predicate> • 251, 301, 372, 401, 402, **411**, 485, 1009, 1218, 1261
- <between predicate part 2> • 211, **411**
- BIGINT • 14, 15, 30, 147, 174, 178, 182, 185, 469, 474, 815, 822, 841, 843, 844, 845, 847, 984, 986, 1076, 1077, 1085, 1092, 1116, 1209, 1230
- BIGINT • 840
- BINARY • 14, 15, 28, 147, 174, 176, 177, 182, 229, 453, 469, 474, 826, 827, 829, 838, 839, 840, 841, 843, 844, 845, 846, 984, 985, 989, 1001, 1031, 1072, 1073, 1079, 1082, 1084, 1088, 1090, 1091, 1094, 1096, 1105, 1107, 1111, 1112, 1113, 1116, 1206
- <binary concatenation> • 29, **268**, 269, 271, 1283
- <binary factor> • **268**, 269, 271
- binary group aggregate functions • 64
- binary large object locator parameter • 811
- binary large object locator variable • 1074, 1083, 1090, 1097, 1102, 1108, 1114
- binary large object string • 826
- <binary large object string type> • **174**, 184, 185, 1202, 1204
- <binary overlay function> • 29, **275**, 279, 280, 286, 287, 1215
- <binary position expression> • 256, **257**, 261
- <binary primary> • **268**, 269, 270
- <binary set function> • 205, **547**, 550, 557, 558, 1174, 1175
- <binary set function type> • **547**, 553, 558, 1222, 1236
- <binary string literal> • 143, 148, 152, **153**, 155, 156, 157, 160, 590, 1202, 1204, 1229, 1255, 1282
- <binary string type> • 173, **174**, 177, 178, 185, 1021, 1202
- binary string types • 14, 15, 224, 1230
- binary strings • iv, 8, 10, 14, 29, 269, 271
- <binary substring function> • 29, **275**, 279, 286, 1283
- <binary trim function> • 29, **275**, 279, 280, 286, 287, 1283
- <binary trim operands> • **275**
- <binary trim source> • **275**, 280
- <binary value expression> • 257, 261, **268**, 269, 271, 275, 279, 280, 286, 414, 415, 479
- <binary value function> • 273, **274**, 275, 279, 280, 286, 287, 1203, 1207



BLOB • 147, 174, 176, 474, 1060, 1062, 1072, 1074, 1080, 1082, 1083, 1084, 1088, 1090, 1095, 1096, 1097, 1101, 1102, 1105, 1106, 1107, 1108, 1112, 1113, 1114

BOOLEAN • 14, 33, 99, 147, 160, 175, 180, 183, 184, 252, 302, 317, 470, 557, 588, 594, 595, 599, 613, 751, 815, 825, 828, 840, 842, 843, 844, 846, 847, 984, 986, 1006, 1076, 1085, 1099, 1105, 1109, 1204, 1287

<boolean factor> • 73, **299**, 301

<boolean literal> • 152, **155**, 160, 590, 1204

<boolean predicand> • 73, 214, **299**, 300, 315, 316, 317, 1168, 1204

<boolean primary> • **299**, 300, 301, 302, 1006, 1204

<boolean term> • 73, **299**

<boolean test> • 73, **299**, 300, 303, 1177

<boolean type> • 173, **175**, 183, 184, 1204

<boolean value expression> • 33, 52, 67, 73, 250, 252, **299**, 300, 301, 302, 315, 316, 317, 444, 451, 601, 661, 1204, 1233

BOTH • 147, 224, 225, 226, 227, 228, 274, 278, 280, 285, 287, 961, 962, 971, 972, 974, 975, 977, 978, 979, 980, 993, 1012, 1036

BPK-set • 72, 73, 74, 76, 77, 78, 79, 80

BPK-sets • 73, 74, 76, 77, 79

<bracketed comment> • **145**, 149, 151, 1217

<bracketed comment contents> • **145**, 149

<bracketed comment introducer> • **145**, 149

<bracketed comment terminator> • **145**

branch • 85, 127, 869, 871, 876, 881, 889, 908, 910, 947, 948, 956, 958, 1149, 1227

branch transaction • 127

*branch transaction already active* • 947, 1149

BREADTH • 146, 393, 394

BUC-set • 72, 73, 74, 76, 77, 78, 79, 80

BUC-sets • 79

BY • 60, 147, 335, 345, 346, 347, 350, 351, 352, 353, 354, 356, 371, 376, 377, 393, 485, 548, 555, 556, 576, 584, 587, 687, 749, 752, 766, 776, 777, 783, 786, 787, 1200, 1218, 1249

## — C —

C • 146, 489, 513, 529, 731, 739, 823, 825, 827, 828, 829, 1021, 1059, 1164

<C array locator variable> • 1079, **1080**, 1083, 1085, 1197

<C array specification> • **1078**, 1079, 1081, 1084

<C BINARY variable> • **1079**, 1082, 1086, 1203

<C BLOB locator variable> • 1079, **1080**, 1083, 1085, 1205

<C BLOB variable> • 1079, **1080**, 1082, 1085, 1205

<C character type> • **1078**, 1081

<C character variable> • **1078**, 1081, 1084, 1242

<C class modifier> • **1078**

<C CLOB locator variable> • 1079, **1080**, 1083, 1085, 1205

<C CLOB variable> • **1079**, 1081, 1082, 1084, 1085, 1205, 1242

<C derived variable> • **1078**

<C host identifier> • 1056, **1078**, 1079, 1080, 1081, 1082, 1084

<C initial value> • 1078, 1079, **1080**, 1084

<C multiset locator variable> • 1079, **1080**, 1083, 1085, 1198

<C NCHAR variable> • **1079**, 1081, 1084

<C NCHAR VARYING variable> • **1079**, 1081, 1084

<C NCLOB variable> • **1079**, 1081, 1082, 1084

<C numeric variable> • **1078**, 1085, 1209

<C REF variable> • 1079, **1080**, 1084, 1085, 1190

<C storage class> • **1078**

<C user-defined type locator variable> • 1079, **1080**, 1083, 1085, 1196

<C user-defined type variable> • **1079**, 1082, 1085, 1199

<C VARBINARY variable> • **1079**, 1082, 1086, 1203

<C VARCHAR variable> • **1079**, 1081, 1084, 1242

<C variable definition> • 1056, **1078**, 1081, 1084

<C variable specification> • **1078**

CALL • 147, 943

<call statement> • 88, 89, 93, 108, 112, 114, 115, 117, 472, 516, 517, 519, 538, 832, **943**, 992, 1011, 1021, 1130, 1162, 1241, 1273

call type data item • 489, 490, 491, 528

CALLED • 147, 677, 680, 698, 717, 720

caller language • 812, 814, 823, 824, 825, 826, 827, 828, 829

candidate basis • 199, 367

candidate key • 57, 80, 372, 373, 581, 629, 1020

candidate original method specification • 703

candidate overriding method specification • 683, 704

candidate routines • 519, 520

CARDINALITY • 147, 257, 327, 444, 446, 448, 985, 986, 996, 1000, 1021

cardinality • 7, 47, 48, 49, 54, 94, 133, 179, 183, 197, 219, 304, 305, 306, 320, 327, 367, 391, 398, 399, 432, 442, 457, 462, 465, 489, 491, 492, 498, 535, 536, 551, 552, 553, 859, 860, 863, 864, 895, 985, 986, 1021, 1031, 1042, 1135, 1231, 1236, 1249

<cardinality expression> • 32, 49, 50, 256, **257**, 260, 262, 266, 1193, 1200, 1232

*cardinality violation* • 398, 399, 863, 882, 1042, 1144

CASCADE • 69, 70, 133, 146, 570, 571, 597, 599, 610, 611, 616, 617, 620, 621, 623, 624, 635, 636, 637, 647, 651, 655, 659, 664, 671, 672, 673, 713, 714, 743, 744, 747, 748, 753, 754, 762, 764, 765, 769, 800, 801, 802, 898, 934, 935, 936, 1165

CASCADED • 57, 60, 147, 625, 626, 627, 631, 668, 924, 931

CASE • 147, 208, 211, 212, 214, 308, 309, 311, 327, 395, 1140, 1168  
 <case abbreviation> • 211, 212, 1007, 1271  
 <case expression> • 188, 189, 211, 212, 214, 230, 372, 463, 1271  
 <case operand> • 211, 212, 214, 1008, 1168  
 <case specification> • 211, 212, 213, 214, 1007  
 CAST • 50, 147, 196, 208, 215, 217, 218, 219, 220, 221, 224, 225, 226, 227, 228, 291, 292, 294, 295, 296, 327, 351, 372, 389, 409, 456, 461, 533, 556, 675, 681, 686, 688, 691, 698, 713, 714, 716, 722, 744, 745, 747, 825, 828, 893, 944, 1025, 1029, 1061, 1064, 1126, 1139, 1269  
 <cast function> • 45, 50, 51, 218, 219, 453, 454, 458, 744, 745, 746, 747  
 <cast operand> • 215, 216, 217, 218, 229, 251, 1007, 1173, 1191, 1206  
 CAST operator • 50  
 <cast specification> • 18, 34, 36, 71, 98, 188, 189, 215, 216, 217, 218, 219, 220, 228, 229, 251, 629, 747, 1007, 1025, 1029, 1271  
 <cast target> • 215, 229, 747, 1007, 1191  
 <cast to distinct> • 674, 675, 676, 677, 689, 1190  
 <cast to distinct identifier> • 675, 677  
 <cast to ref> • 674, 675, 676, 677, 678, 689, 1190  
 <cast to ref identifier> • 675, 678  
 <cast to source> • 674, 675, 676, 677, 689, 1190  
 <cast to source identifier> • 675, 677  
 <cast to type> • 674, 675, 676, 677, 678, 689, 1190  
 <cast to type identifier> • 675, 678  
 CATALOG • 146, 974, 975  
 <catalog name> • 81, 161, 167, 168, 169, 171, 193, 286, 568, 570, 806, 974, 975, 1136, 1137, 1138, 1139, 1178, 1229, 1238  
 <catalog name characteristic> • 974, 1010  
 CATALOG\_NAME • 146, 1126, 1136, 1137, 1138  
 CEIL • 147, 258  
 CEILING • 147, 258  
 <ceiling function> • 32, 256, 258, 260, 264, 267, 1222, 1232  
 CHAIN • 127, 146, 955, 956, 957, 958, 959, 1213  
 CHAR • 147, 173, 174, 176, 474, 815, 821, 822, 824, 1073, 1074, 1075, 1076, 1105, 1106, 1107, 1109, 1111  
 <char length expression> • 257, 261, 262, 1005, 1259, 1283  
 <char length units> • 174, 176, 177, 180, 185, 256, 257, 259, 261, 262, 265, 267, 273, 274, 276, 277, 279, 281, 284, 287, 838, 1073, 1081, 1084, 1089, 1096, 1106, 1109, 1113, 1115, 1185, 1208  
 CHAR\_LENGTH • 147, 257, 261, 278  
 CHARACTER • 14, 15, 16, 18, 98, 147, 157, 173, 174, 176, 177, 182, 186, 228, 229, 266, 419, 425, 454, 468,

474, 539, 567, 571, 585, 648, 650, 651, 779, 821, 822, 823, 824, 825, 826, 827, 828, 829, 838, 839, 841, 842, 843, 844, 845, 846, 984, 985, 990, 1001, 1005, 1006, 1009, 1010, 1011, 1031, 1071, 1072, 1073, 1076, 1078, 1079, 1081, 1084, 1087, 1088, 1089, 1092, 1094, 1095, 1096, 1098, 1100, 1101, 1105, 1106, 1109, 1111, 1112, 1113, 1115, 1139, 1173, 1206, 1207, 1225, 1230  
 character • 5  
 character comparison operation • 19  
 character encoding form • 5  
 <character enumeration> • 421, 423, 424  
 <character enumeration exclude> • 421, 424  
 <character enumeration include> • 421, 424  
 <character factor> • 268, 269, 270, 283, 1233  
 <character large object length> • 173, 174, 177, 274, 279, 1072, 1079, 1081, 1087, 1088, 1089, 1094, 1096, 1100, 1105, 1106, 1111, 1113  
 character large object locator parameter • 812  
 character large object locator variable • 1075, 1083, 1090, 1097, 1103, 1108, 1114  
 <character large object type> • 18, 173, 184, 269, 270, 277, 455, 460, 464, 532, 533, 680, 681, 698, 699, 723, 811, 812, 1204, 1282  
 <character length> • 157, 173, 174, 176, 274, 279, 1071, 1073, 1078, 1081, 1084, 1087, 1089, 1094, 1096, 1100, 1101, 1105, 1106, 1109, 1111, 1113  
 <character like predicate> • 414, 415, 419, 1009, 1169  
 <character like predicate part 2> • 211, 414  
 character not in repertoire • 180, 1145  
 <character overlay function> • 21, 29, 273, 274, 275, 278, 281, 287, 1215  
 <character pattern> • 414, 419, 478, 1009, 1169  
 <character position expression> • 256, 257, 258, 261  
 <character primary> • 268, 269, 270  
 <character representation> • 152, 155, 156, 157, 158, 160, 1071, 1076, 1080, 1084, 1087, 1091, 1111, 1115, 1169, 1259  
 character set collation • 19, 192  
 <character set definition> • 105, 166, 539, 567, 648, 649, 831, 1130, 1175  
 <character set name> • 156, 162, 166, 168, 169, 171, 176, 539, 571, 639, 648, 649, 650, 656, 657, 772, 779, 780, 789, 821, 977, 1175, 1229, 1230  
 <character set name characteristic> • 977, 1010  
 <character set source> • 648  
 <character set specification> • 26, 150, 152, 156, 157, 164, 173, 177, 186, 539, 540, 567, 569, 585, 638, 648, 649, 652, 656, 690, 772, 801, 805, 806, 807, 810, 838, 980, 985, 1056, 1071, 1072, 1073, 1076, 1078, 1079, 1081, 1084, 1087, 1088, 1089, 1092, 1094, 1096, 1098, 1100, 1105, 1106, 1109, 1111, 1113, 1115, 1175, 1237, 1241, 1242

- <character set specification list> • **805**, 807, 980, 981
- <character specifier> • **420**, 421, 422, 423
- <character string literal> • **144**, 148, 150, **152**, 155, 156, 157, 158, 160, 162, 164, 589, 590, 977, 1169
- <character string type> • 18, 19, 34, **173**, 176, 177, 182, 252, 275, 276, 278, 377, 484, 498, 550, 612, 626, 638, 678, 690, 821, 825, 1020, 1126, 1230, 1238, 1250, 1259
- character string types • 14, 1230, 1260, 1269, 1282
- character strings • 8, 10, 14, 19, 22, 23, 29, 98, 269, 270, 279, 408, 416, 424, 476, 487, 828, 1004, 1084, 1139
- <character substring function> • 20, 29, **273**, 275, 281, 1259, 1283
- <character transliteration> • 22, **273**, 275, 277, 281, 284, 287, 1180, 1233
- <character value expression> • 256, 257, 261, **268**, 269, 270, 271, 273, 274, 275, 276, 277, 278, 279, 281, 282, 283, 284, 285, 414, 419, 420, 421, 422, 423, 425, 1173, 1207, 1233
- <character value expression 1> • **257**, 258, 261
- <character value expression 2> • **257**, 258, 261
- <character value function> • **273**, 275, 280, 281
- CHARACTER\_LENGTH • 8, 147, 257, 265, 281, 284, 476
- CHARACTER\_SET\_CATALOG • 146, 984, 985, 996, 1000, 1001, 1002, 1020, 1241
- CHARACTER\_SET\_NAME • 146, 984, 985, 996, 1000, 1001, 1002, 1020, 1241
- CHARACTER\_SET\_SCHEMA • 146, 984, 985, 996, 1000, 1001, 1002, 1020, 1241
- CHARACTERISTICS • 146, 969
- CHARACTERS • 146, 174, 176, 177, 180, 259, 261, 276, 277, 840, 842, 844
- CHECK • 57, 60, 147, 586, 601, 602, 625, 626, 627, 631, 633, 647, 661, 668, 877, 882, 884, 890, 911, 924, 931, 1137, 1178, 1181, 1209
- <check constraint definition> • 203, 229, 332, 584, 586, 593, 594, **601**, 602, 638, 1178, 1179, 1266
- <circumflex> • 23, 140, **141**, 421, 422, 424
- CLASS\_ORIGIN • 146, 1126, 1136, 1152, 1243
- CLOB • 147, 173, 176, 474, 1060, 1062, 1072, 1073, 1074, 1079, 1080, 1081, 1083, 1084, 1087, 1088, 1089, 1090, 1094, 1095, 1096, 1097, 1100, 1101, 1102, 1103, 1105, 1107, 1108, 1111, 1112, 1113, 1114
- CLOSE • 147, 861, 1043
- close call • 490, 491
- <close statement> • 101, 106, 110, 113, 116, 117, 832, **861**, 1130, 1264, 1265
- CLU • 265
- COALESCE • 8, 76, 147, 211, 212, 327, 337, 342
- COBOL • 99, 146, 489, 513, 529, 731, 739, 741, 808, 823, 825, 826, 827, 828, 829, 1059, 1088, 1089, 1090, 1091, 1092, 1153, 1159, 1163, 1164, 1242
- <COBOL array locator variable> • 1087, **1088**, 1091, 1092, 1197
- <COBOL binary integer> • **1088**, 1092
- <COBOL BINARY variable> • 1087, **1088**, 1090, 1093, 1203
- <COBOL BLOB locator variable> • 1087, **1088**, 1090, 1092, 1205
- <COBOL BLOB variable> • 1087, **1088**, 1090, 1092, 1205
- <COBOL character type> • **1087**, 1089, 1092
- <COBOL CLOB locator variable> • 1087, **1088**, 1090, 1092, 1205
- <COBOL CLOB variable> • **1087**, 1089, 1092, 1205
- <COBOL derived type specification> • **1087**
- <COBOL host identifier> • 1056, **1087**, 1089, 1090, 1091
- <COBOL integer type> • 1087, **1088**
- <COBOL multiset locator variable> • 1087, **1088**, 1091, 1092, 1198
- <COBOL national character type> • **1087**, 1088, 1089
- <COBOL NCLOB variable> • **1087**, 1089
- <COBOL nines> • **1088**
- <COBOL nines specification> • **1088**
- <COBOL numeric type> • 1087, **1088**, 1092
- <COBOL REF variable> • 1087, **1088**, 1091, 1092, 1190
- <COBOL type specification> • **1087**
- <COBOL user-defined type locator variable> • 1087, **1088**, 1091, 1092, 1196
- <COBOL user-defined type variable> • 1087, **1088**, 1090, 1092, 1199
- <COBOL variable definition> • 1056, **1087**, 1089, 1091, 1242
- code point • **5**
- code unit • **6**
- COLLATE • 147, 544
- <collate clause> • 62, 173, 176, 215, 268, 269, 345, 346, 356, 357, 359, **544**, 584, 585, 588, 638, 640, 648, 649, 690, 691, 1179
- COLLATION • 146, 191, 571, 652, 654, 655, 772, 779, 802, 805, 980
- collation • **5**, 18, 19, 24, 25, 26, 27, 28, 81, 84, 119, 121, 130, 131, 150, 157, 176, 182, 191, 192, 193, 200, 216, 261, 268, 269, 275, 276, 277, 278, 346, 357, 377, 408, 415, 421, 425, 463, 487, 540, 544, 571, 588, 626, 640, 648, 650, 652, 653, 654, 655, 659, 691, 772, 788, 789, 792, 793, 794, 795, 797, 798, 800, 801, 802, 807, 980, 981, 1020, 1148, 1154, 1179, 1226, 1229, 1237, 1249, 1278
- <collation definition> • 105, 166, 567, **652**, 653, 789, 831, 1130, 1179, 1237, 1249
- <collation name> • 62, **162**, 166, 169, 171, 346, 357, 544, 571, 639, 648, 649, 652, 654, 772, 779, 780, 789, 805, 807, 980, 1179
- <collation specification> • **980**

COLLATION\_CATALOG • 146, 996, 999, 1020, 1241

COLLATION\_NAME • 146, 996, 999, 1020, 1241

COLLATION\_SCHEMA • 146, 996, 999, 1020, 1241

COLLECT • 147, 547, 549, 558, 1201

collection • iv, 7, 8, 9, 16, 32, 47, 48, 49, 53, 61, 64, 66, 97, 125, 133, 183, 184, 197, 215, 217, 218, 316, 339, 340, 341, 360, 454, 458, 470, 498, 547, 549, 560, 562, 589, 590, 591, 693, 695, 712, 721, 775, 799, 800, 883, 894, 983, 1202, 1281

<collection derived table> • 76, 325, **326**, 329, 331, 333, 334, 1193, 1201, 1202

<collection type> • 47, 48, 173, **175**, 180, 183, 184, 1202

collection type descriptor • 48

collection types • 7, 13, 16, 47, 404

<collection value constructor> • **188**, 189

<collection value expression> • **250**, 252, 257, 260, 262, 326, 334, 463, 1202

<colon> • **140**, 154, 162, 421, 422, 424, 511, 512, 1056, 1071, 1105

COLUMN • 147, 604, 606, 616, 617, 801, 802, 1165

column • 7, 44, 46, 47, 51, 52, 53, 54, 55, 56, 58, 60, 61, 62, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 85, 103, 119, 120, 121, 134, 135, 136, 192, 193, 198, 199, 200, 201, 202, 203, 205, 215, 241, 243, 251, 252, 299, 300, 306, 307, 313, 317, 320, 322, 323, 330, 331, 332, 336, 337, 338, 339, 340, 341, 342, 343, 344, 346, 350, 352, 353, 354, 355, 357, 358, 359, 365, 367, 368, 369, 370, 371, 372, 373, 374, 375, 377, 381, 382, 383, 384, 385, 387, 389, 398, 399, 408, 419, 433, 436, 439, 463, 478, 479, 481, 485, 487, 490, 491, 529, 552, 553, 555, 556, 557, 558, 561, 562, 564, 565, 566, 575, 576, 577, 578, 579, 580, 581, 582, 584, 585, 586, 587, 588, 589, 590, 591, 592, 595, 598, 599, 600, 603, 604, 605, 606, 607, 608, 609, 610, 611, 612, 614, 615, 616, 617, 618, 619, 621, 623, 626, 627, 629, 630, 631, 632, 635, 636, 639, 643, 644, 645, 646, 650, 655, 662, 663, 664, 668, 669, 670, 693, 695, 712, 731, 732, 741, 771, 772, 773, 774, 777, 780, 781, 787, 788, 789, 790, 791, 792, 793, 794, 795, 796, 797, 798, 799, 800, 801, 802, 814, 823, 824, 858, 863, 874, 875, 876, 877, 879, 880, 883, 884, 885, 886, 892, 893, 894, 895, 896, 897, 913, 917, 919, 920, 922, 923, 925, 926, 928, 930, 933, 934, 935, 936, 937, 997, 999, 1008, 1018, 1019, 1020, 1028, 1046, 1047, 1053, 1074, 1082, 1090, 1097, 1102, 1107, 1114, 1137, 1144, 1169, 1171, 1172, 1174, 1175, 1177, 1181, 1183, 1189, 1212, 1214, 1222, 1226, 1227, 1233, 1237, 1246, 1247, 1248, 1250, 1257, 1261, 1262, 1263, 1277, 1279, 1285, 1286

<column constraint> • 576, **584**, 586, 1265, 1266

<column constraint definition> • 574, 576, 578, 579, **584**, 586, 587

<column default option> • **574**, 575

<column definition> • 71, 177, 184, 569, 573, 575, 577, 578, 579, 580, 582, **584**, 585, 587, 588, 589, 596, 604, 605, 897, 898, 1179, 1191, 1220, 1266

column list • 135

<column name> • 58, **162**, 169, 198, 199, 200, 202, 326, 327, 328, 330, 331, 337, 366, 368, 370, 371, 375, 377, 382, 383, 386, 393, 394, 395, 396, 574, 577, 578, 579, 584, 585, 595, 596, 598, 604, 605, 606, 607, 608, 609, 610, 612, 615, 616, 625, 626, 629, 630, 631, 667, 716, 718, 772, 780, 781, 788, 789, 790, 791, 794, 795, 798, 801, 802, 855, 874, 875, 876, 879, 880, 884, 887, 892, 1047, 1053, 1063, 1137, 1214, 1220, 1248

<column name list> • 120, **326**, 336, 366, 376, 377, 393, 394, 574, 577, 595, 597, 625, 666, 779, 853, 854, 855, 873, 887, 1047, 1053, 1255, 1266

<column option list> • **574**, 579, 582, 1191

<column options> • 573, **574**, 578, 579

column privilege descriptor • 120

<column reference> • 61, 63, 136, 188, 189, 190, 191, 192, 193, 200, **202**, 203, 205, 206, 252, 300, 335, 344, 345, 350, 351, 354, 356, 357, 369, 370, 371, 374, 375, 385, 386, 520, 535, 576, 585, 617, 791, 792, 793, 794, 795, 796, 798, 857, 862, 886, 1053, 1174, 1184, 1214, 1218, 1246, 1264, 1265, 1272

COLUMN\_NAME • 146, 1126, 1137

<comma> • 139, **140**, 174, 175, 193, 205, 211, 257, 258, 274, 306, 312, 315, 320, 323, 326, 335, 345, 347, 348, 349, 350, 356, 366, 370, 376, 393, 412, 420, 449, 509, 515, 516, 541, 547, 548, 559, 573, 625, 674, 675, 715, 716, 717, 755, 758, 761, 776, 779, 783, 786, 805, 811, 857, 862, 878, 892, 899, 900, 949, 951, 969, 996, 999, 1023, 1027, 1071, 1078, 1079, 1080, 1094, 1100, 1105, 1111, 1125

COMMAND\_FUNCTION • 146, 1125, 1129, 1139

COMMAND\_FUNCTION\_CODE • 146, 1125, 1129, 1243

<comment> • **145**, 149

<comment character> • **145**

COMMIT • 57, 125, 126, 147, 573, 579, 581, 598, 601, 897, 955

<commit statement> • 68, 94, 101, 104, 108, 110, 112, 123, 125, 127, 128, 808, 832, **955**, 956, 966, 1120, 1130, 1213, 1251, 1266

COMMITTED • 124, 125, 127, 146, 948, 949

common column name • 336, 337

<common sequence generator option> • **766**

<common sequence generator options> • 505, 584, 586, **766**, 767

<common value expression> • 73, 214, **250**, 252, 299, 315, 316, 317, 414, 419, 421, 426, 434, 444, 446, 448, 449, 1168, 1169

<comp op> • 8, 73, **403**, 404, 405, 406, 407, 408, 410, 430, 431, 560, 1249, 1268

comparable • 7

- comparison category • 44, 45, 404, 406, 407, 750, 1234
- comparison form • 17, 44, 404, 442, 750
- comparison function • 44, 45, 404, 406, 407
- <comparison predicate> • 8, 23, 29, 73, 191, 251, 301, 372, 401, 402, **403**, 404, 430, 431, 478, 483, 485, 560, 1009, 1268, 1269, 1270
- <comparison predicate part 2> • 211, **403**
- comparison type • 18, 44, 45, 433, 753, 754
- compatible • 9, 16, 43, 88, 216, 218, 463, 498, 535, 550, 628, 680, 681, 682, 699, 700, 701, 705, 706, 719, 756, 759, 760, 1026, 1030, 1255
- compilation unit • **6**
- completion condition • 95, 96, 103, 123, 222, 223, 224, 284, 455, 456, 457, 490, 536, 537, 552, 553, 592, 777, 836, 837, 858, 863, 870, 872, 877, 882, 884, 890, 906, 909, 911, 938, 948, 967, 997, 1034, 1041, 1042, 1122, 1123, 1125, 1143, 1144
- <computational operation> • **547**, 557, 558, 1174, 1201, 1202, 1204, 1208, 1222, 1263
- <concatenation> • 19, **268**, 269, 270, 271, 340, 347, 348, 497, 883, 1259, 1283
- <concatenation operator> • 20, 144, **145**, 268, 304, 1006
- concurrent SQL-transactions • 124, 125, 126
- CONDITION • 147, 1125, 1126, 1140, 1255
- <condition> • **1067**, 1068
- <condition action> • **1067**, 1069
- condition area limit • 98, 124, 127, 131, 524, 537, 948, 956, 958, 962, 970, 1245
- condition areas • 97, 98, 537, 1129, 1136, 1139
- <condition information> • **1125**, 1136, 1140
- <condition information item> • **1125**, 1126, 1136, 1140
- <condition information item name> • **1125**, 1126, 1127, 1139, 1140, 1141, 1167
- condition information items • 97
- <condition number> • 1125, **1126**, 1127, 1136, 1140, 1251
- CONDITION\_NUMBER • 146, 1126, 1136, 1139, 1140
- conflicting method specification • 681, 699, 700
- conflicting overriding method specification • 705
- CONNECT • 147, 835, 961, 1121
- <connect statement> • 108, 118, 127, 128, 129, 137, 832, 835, **961**, 962, 963, 1120, 1130, 1182, 1228, 1240, 1250
- CONNECTION • 146, 835, 964, 1121
- connection does not exist* • 835, 964, 966, 1121, 1144
- connection exception* • 128, 835, 962, 964, 966, 1121, 1144
- connection failure* • 964, 1144
- <connection name> • 127, **163**, 168, 169, 171, 961, 962, 964, 966, 1139, 1182
- connection name in use* • 962, 1144
- <connection object> • **964**, 965, 966
- <connection target> • **961**
- <connection user name> • 129, **163**, 168, 961, 962, 963, 1240
- CONNECTION\_NAME • 146, 1126, 1139
- considered as executed • 135, 913, 917, 920, 923, 926, 930, 940
- constituent • 15, 16, 94, 251, 338, 384, 598
- constituent rows • 94
- CONSTRAINT • 147, 545, 611, 619, 620, 624, 645, 664, 754, 801, 1067, 1068, 1069, 1070, 1163
- <constraint characteristics> • **545**, 546, 584, 586, 593, 638, 647, 661, 662, 1181
- <constraint check time> • **545**
- constraint mode • 67, 68, 124, 127, 131, 524, 537, 545, 588, 593, 639, 662, 836, 951, 956, 958, 1121
- <constraint name> • **162**, 167, 169, 171, 545, 571, 593, 611, 619, 620, 638, 639, 645, 647, 661, 662, 663, 664, 748, 754, 801, 951, 1067, 1069, 1070, 1176, 1249
- <constraint name definition> • **545**, 546, 584, 586, 593, 638, 639, 647, 1176, 1181, 1249
- <constraint name list> • **951**
- CONSTRAINT\_CATALOG • 146, 1126, 1136, 1138
- CONSTRAINT\_NAME • 146, 1126, 1136, 1138
- CONSTRAINT\_SCHEMA • 146, 1126, 1136, 1138
- CONSTRAINTS • 146, 951, 955
- constructed types • 13, 16, 50, 477
- CONSTRUCTOR • 40, 42, 146, 541, 542, 543, 676, 679, 680, 681, 682, 683, 685, 697, 699, 700, 701, 703, 704, 706, 708, 709, 710, 711, 716, 718, 719, 735
- constructor function • **7**, 43, 240, 522, 687
- <constructor method selection> • **235**, 236, 517, 518, 519
- containing schema • 82, 203, 217, 218, 230, 244, 246, 280, 332, 517, 540, 544, 650, 654, 732, 743
- containing SQL • 130
- containing SQL not permitted* • 530
- containing SQL not permitted* • 1147
- CONTAINS • 146, 677, 698, 717, 725, 732, 736, 742, 751, 1238
- contaminated • 799, 802
- <contextually typed row value constructor> • **315**, 316, 317, 318, 319, 321, 874, 1178, 1208, 1247
- <contextually typed row value constructor element> • **315**, 316, 317, 874
- <contextually typed row value constructor element list> • **315**, 317, 321, 1178
- <contextually typed row value expression> • 231, **318**, 319, 320, 874, 875, 892, 1008, 1010, 1264
- <contextually typed row value expression list> • **320**, 321, 1178
- <contextually typed table value constructor> • 231, 317, **320**, 321, 873, 874, 875, 877, 885, 1008, 1178, 1189, 1247, 1264

<contextually typed value specification> • **196**, 315, 316, 875, 878, 879, 880, 892, 894, 1271  
CONTINUE • 146, 1067  
control character • **6**  
CONVERT • 147, 273  
Coordinated Universal Time • **6**  
CORR • 65, 147, 548, 553  
<correlation name> • 75, **162**, 168, 169, 199, 325, 326, 327, 329, 367, 368, 394, 667, 866, 868, 869, 878, 879, 881, 883, 884, 886, 888, 889, 1135, 1183, 1261, 1262  
CORRESPONDING • 147, 376, 382, 391, 616, 633, 1169  
<corresponding column list> • **376**, 382  
corresponding fields • 403  
corresponding join columns • 75, 336, 337, 339, 342, 343  
<corresponding spec> • **376**  
corresponding subfields • 1004  
COUNT • 64, 147, 208, 372, 547, 549, 551, 552, 557, 996, 997, 1018, 1024, 1028, 1134, 1135, 1174, 1231, 1236  
counterpart • 72, 73, 74, 77, 78, 79, 373, 562, 564, 565, 592  
COVAR\_POP • 65, 147, 547, 553  
COVAR\_SAMP • 65, 147, 547, 553  
CREATE • 147, 567, 573, 625, 626, 638, 648, 652, 656, 661, 666, 674, 686, 687, 688, 715, 745, 749, 751, 755, 766, 782  
created by • 81  
CROSS • 74, 147, 335, 338  
<cross join> • **335**, 339, 341, 343, 1172  
CUBE • 63, 73, 78, 147, 345, 348  
<cube list> • **345**, 346, 347, 348, 349, 353, 1217  
CUME\_DIST • 63, 64, 66, 147, 207, 208, 210, 556, 1221, 1227  
CURRENT • 147, 208, 357, 358, 361, 363, 364, 365, 866, 886, 966, 1044, 1046, 1048, 1052  
current • 36, 61, 67, 68, 82, 83, 90, 91, 94, 95, 101, 102, 110, 111, 112, 118, 119, 121, 122, 123, 128, 129, 130, 131, 132, 158, 160, 164, 165, 166, 167, 180, 192, 193, 203, 217, 218, 225, 231, 240, 244, 246, 280, 289, 292, 332, 361, 362, 363, 364, 365, 487, 504, 517, 518, 519, 524, 526, 531, 537, 538, 540, 544, 549, 623, 630, 733, 735, 736, 772, 774, 804, 813, 834, 835, 836, 856, 858, 859, 860, 861, 863, 864, 867, 869, 871, 875, 876, 880, 881, 887, 889, 898, 905, 906, 908, 909, 910, 911, 919, 934, 935, 941, 945, 947, 948, 951, 953, 954, 955, 956, 957, 958, 961, 962, 963, 964, 965, 966, 967, 969, 971, 972, 973, 974, 975, 977, 978, 979, 980, 981, 1025, 1028, 1029, 1037, 1038, 1041, 1044, 1046, 1049, 1050, 1053, 1065, 1088, 1120, 1121, 1229, 1235, 1240, 1241, 1249, 1250  
current authorization identifier • 118, 119, 617, 636, 647, 671, 672, 673, 774, 802, 1166  
current base value • 82, 83, 504, 506, 508

<current collation specification> • 190, **191**, 192, 193, 195, 1179  
<current date value function> • **291**, 1269  
<current local time value function> • **291**, 292, 1177, 1269  
<current local timestamp value function> • **291**, 292, 1177, 1269  
current role name • 118, 119, 122, 193, 804, 897, 971, 972, 1240  
current savepoint level • 123  
current SQL-connection • 128  
current SQL-session • 129, 524, 956, 958  
<current time value function> • **291**, 292, 1173  
<current timestamp value function> • **291**, 292, 1173  
current user identifier • 118, 119, 122, 193, 526, 772, 773, 774, 804, 897, 971, 972  
CURRENT\_CATALOG • 129, 147, 190, 191, 193, 194, 251, 589, 590, 591, 592, 902, 1123, 1182, 1231, 1255  
CURRENT\_DATE • 126, 147, 226, 227, 291, 292, 301, 602, 662, 1179, 1247  
CURRENT\_DEFAULT\_TRANSFORM\_GROUP • 130, 147, 190, 192, 194, 902, 1123, 1198  
CURRENT\_PATH • 95, 130, 147, 190, 191, 193, 194, 251, 589, 590, 591, 592, 902, 1123, 1192, 1231  
CURRENT\_ROLE • 118, 147, 190, 191, 193, 194, 251, 372, 589, 590, 591, 592, 779, 804, 902, 1123, 1216  
CURRENT\_SCHEMA • 130, 147, 190, 191, 193, 194, 251, 589, 590, 591, 592, 902, 1123, 1182, 1231, 1255  
CURRENT\_TIME • 147, 291, 292, 1247  
CURRENT\_TIMESTAMP • 147, 291, 292, 301, 602, 662, 736, 1179, 1247  
CURRENT\_TRANSFORM\_GROUP\_FOR\_TYPE • 130, 147, 190, 192, 194, 902, 1123, 1199  
CURRENT\_USER • 118, 126, 147, 190, 191, 193, 194, 251, 372, 589, 590, 591, 592, 779, 804, 902, 1123, 1170, 1231  
CURSOR • 147, 851, 854, 1017, 1036  
<cursor attribute> • **1014**  
<cursor attributes> • 102, **1014**  
cursor declaration descriptor • 57, 99, 100, 101, 102, 526, 813, 849, 851, 856, 857, 858, 861, 866, 867, 886, 887, 901, 907, 1015, 1035, 1037, 1038, 1039  
<cursor holdability> • 102, **851**, 852, 854, 902, 1014, 1219, 1264, 1265  
cursor instance descriptor • 100, 101, 131, 526, 537, 813, 856, 858, 861, 867, 887, 1015, 1037, 1038  
<cursor intent> • 853, **1036**  
<cursor name> • 99, 100, **162**, 163, 165, 169, 170, 171, 807, 849, 850, 854, 856, 857, 861, 866, 886, 1035, 1039, 1041, 1043, 1044, 1046, 1050, 1058, 1220, 1271  
cursor operation conflict • 870, 872, 882, 890, 909, 911, 938, 1136, 1150

<cursor properties> • 102, 849, 850, **851**, 853, 854, 1035, 1036, 1037, 1184, 1219  
 <cursor returnability> • 102, **851**, 852, 854, 902, 1014, 1218  
 <cursor scrollability> • 102, 850, **851**, 852, 854, 902, 1014, 1174, 1184  
 <cursor sensitivity> • 102, **851**, 852, 854, 902, 1014, 1183, 1213  
*cursor sensitivity exception* • 869, 871, 876, 881, 889, 903, 908, 910, 1144  
 <cursor specification> • 9, 57, 86, 94, 100, 101, 102, 103, 230, 391, 794, 795, 797, 798, 799, 849, 850, **853**, 855, 857, 886, 887, 901, 904, 1004, 1015, 1035, 1036, 1037, 1039, 1040, 1046, 1047, 1048, 1052, 1053, 1123, 1184, 1185, 1209, 1246, 1255  
*cursor specification cannot be executed* • 1032, 1147  
 CURSOR\_NAME • 146, 1126, 1136, 1137  
 CYCLE • 83, 147, 393, 504, 506, 767  
 <cycle clause> • **393**, 394, 395, 396  
 <cycle column> • **393**  
 <cycle column list> • **393**, 394  
 <cycle mark column> • **393**, 394  
 <cycle mark value> • **393**, 394  
 cycle option • 60

## — D —

DATA • 56, 146, 574, 582, 612, 717, 720, 725, 732, 736, 742, 829, 996, 997, 998, 999, 1000, 1019, 1024, 1025, 1030, 1074, 1237, 1250  
*data exception* • 180, 193, 194, 219, 220, 221, 222, 223, 224, 225, 226, 227, 228, 229, 248, 255, 262, 263, 264, 265, 270, 271, 281, 282, 285, 286, 287, 290, 297, 305, 307, 320, 333, 361, 362, 363, 364, 416, 417, 422, 454, 455, 456, 457, 459, 460, 461, 462, 490, 493, 496, 504, 522, 535, 536, 551, 552, 555, 557, 691, 813, 825, 859, 860, 864, 865, 895, 896, 973, 997, 998, 1000, 1139, 1144, 1231, 1233  
 <data type> • **6**, 19, 46, 47, 85, 92, 93, 95, 157, 168, **173**, 175, 179, 180, 182, 183, 184, 186, 215, 216, 229, 235, 236, 316, 331, 338, 372, 373, 490, 491, 519, 532, 534, 541, 542, 578, 582, 584, 585, 587, 603, 612, 630, 677, 681, 684, 690, 691, 698, 705, 715, 716, 717, 719, 721, 722, 728, 730, 732, 734, 736, 737, 738, 741, 745, 766, 767, 792, 793, 794, 795, 796, 798, 807, 811, 813, 814, 822, 823, 834, 1007, 1059, 1189, 1191, 1194, 1196, 1197, 1198, 1201, 1202, 1205, 1230, 1266, 1282  
 <data type list> • **541**, 542, 708, 709  
 <data type or domain name> • 71, **584**, 585  
*data type transform function violation* • 1026, 1030, 1147  
 DATE • 14, 15, 34, 35, 37, 98, 147, 153, 157, 175, 179, 180, 183, 224, 288, 291, 470, 474, 613, 840, 842, 843, 844, 846, 847, 1001, 1011  
 date • **6**

<date literal> • **153**, 157, 1268  
 <date string> • 144, **153**  
 <date value> • **154**, 159  
 datetime component • 158, 159  
 <datetime factor> • 251, **288**  
*datetime field overflow* • 290, 361, 456, 461, 1145  
 <datetime literal> • 152, **153**, 158, 159, 160, 590, 1268, 1269  
 <datetime primary> • 251, **288**, 289  
 <datetime term> • 251, **288**, 289, 290, 293, 294, 296  
 <datetime type> • 173, **175**, 180, 181, 183, 984, 1021, 1231, 1268, 1269  
 datetime types • 15, 1269  
 <datetime value> • 154, **155**, 158  
 <datetime value expression> • 250, 251, 252, 257, 259, 260, 261, **288**, 289, 290, 293, 294, 296, 1166, 1269  
 <datetime value function> • 251, 288, 289, **291**, 292, 589, 590, 591, 902, 1269  
 datetime with time zone • 35, 203, 251, 252, 289, 301, 374, 384, 456, 461, 553, 1247  
 datetime without time zone • 35, 251, 301, 338, 384, 598  
 DATETIME\_INTERVAL\_CODE • 146, 984, 996, 1000, 1001, 1002, 1021  
 DATETIME\_INTERVAL\_PRECISION • 146, 984, 996, 1000, 1001, 1002, 1021  
 datetimes • iv, 15, 32, 39, 409, 439, 1233  
 DAY • 35, 147, 158, 181, 289, 361, 464, 470, 509, 510, 1002, 1006, 1011  
 day-time • 34, 37, 39, 183, 294, 470, 613  
 <day-time interval> • **154**, 179  
 <day-time literal> • **154**, 158  
 <days value> • **154**, 158  
 DEALLOCATE • 147, 531, 808, 995, 1012, 1015, 1034, 1120  
 <deallocate descriptor statement> • 86, 109, 111, 808, 833, **995**, 1120, 1130, 1160  
 <deallocate prepared statement> • 85, 86, 100, 101, 109, 111, 833, **1015**, 1016, 1130, 1162  
 DEC • 30, 147, 174, 176, 178, 182, 825, 828, 1100, 1101, 1112, 1230  
 DECIMAL • 14, 15, 30, 147, 174, 176, 178, 182, 469, 474, 840, 841, 843, 844, 845, 847, 984, 985, 1001, 1101, 1112, 1116, 1230  
 DECLARE • 147, 849, 897, 1035, 1055, 1056, 1064, 1111, 1115  
 <declare cursor> • 10, 84, 99, 100, 102, 106, 112, 114, 116, 117, 165, 805, 807, **849**, 850, 853, 854, 856, 1035, 1055, 1058, 1060, 1064, 1065, 1184, 1251, 1264, 1265  
 declared cursor • 99, 100, 101, 102, 537, 813  
 declared dynamic cursor • 86, 99, 100, 101, 102, 170, 901, 1015, 1035, 1039, 1041, 1043, 1044, 1046, 1050, 1255

- declared type • 7
- declared type collation • 19, 176, 463, 487
- DEFAULT • 60, 69, 70, 133, 147, 196, 393, 567, 576, 584, 587, 589, 597, 599, 608, 643, 835, 873, 874, 877, 911, 934, 935, 936, 937, 961, 962, 964, 966, 979, 1121, 1168, 1240
- <default clause> • 51, 574, 578, 579, 584, 587, **589**, 590, 592, 607, 638, 639, 642, 646, 690, 874, 1266, 1269
- <default option> • 51, 52, 196, 576, 578, 579, 587, **589**, 590, 591, 592, 1170, 1182, 1192, 1216
- <default specification> • **196**, 197, 316, 874, 877, 879, 880, 884, 893, 1271
- default SQL-connection • 128
- default SQL-session • 129
- default value too long for information schema* • 592, 1151
- DEFAULTS • 146, 574, 575, 576
- DEFERRABLE • 146, 545, 546, 586, 593, 638, 661, 1181
- deferrable • 52, 67, 545, 588, 593, 639, 662, 951
- DEFERRED • 146, 545, 951
- deferred • 67, 68, 69, 70, 122, 123, 127, 545, 593, 639, 951, 1228
- DEFINED • 146, 526, 717, 725, 737, 984, 986, 1236
- DEFINER • 146, 526, 717, 725, 734, 735, 737, 797, 1012
- DEGREE • 146, 984, 986, 996, 1000, 1001, 1021
- degree • 46, 54, 73, 94, 124, 183, 186, 217, 306, 313, 316, 320, 328, 330, 367, 368, 377, 382, 389, 398, 399, 403, 405, 428, 429, 441, 454, 457, 458, 462, 464, 498, 528, 575, 577, 580, 587, 605, 617, 626, 628, 686, 687, 694, 696, 726, 857, 858, 859, 862, 863, 875, 893, 905, 984, 986, 1008, 1018, 1021, 1027
- DELETE • 57, 70, 120, 121, 134, 135, 136, 147, 561, 564, 565, 566, 573, 579, 581, 582, 597, 598, 601, 666, 667, 668, 670, 673, 779, 787, 795, 798, 836, 866, 867, 868, 869, 897, 898, 913, 917, 940, 955, 1044, 1048, 1049
- delete INSTEAD OF trigger • 133
- <delete rule> • 69, 70, **597**, 599, 600, 933, 1167
- <delete statement: positioned> • 60, 94, 102, 103, 107, 110, 114, 116, 117, 795, 798, 832, **866**, 867, 870, 872, 882, 890, 909, 911, 938, 1130, 1264, 1265
- <delete statement: searched> • 60, 107, 110, 114, 115, 117, 794, 795, 796, 798, 799, 832, **868**, 870, 909, 911, 955, 1003, 1119, 1130, 1134, 1135, 1183, 1209, 1251, 1264
- delete trigger • 133
- <delimited identifier> • **144**, 148, 149, 150, 151, 161, 193, 1171, 1225, 1229, 1260
- <delimited identifier body> • **144**, 149, 150, 151, 1171
- <delimited identifier part> • **144**, 149, 151, 1171
- <delimiter token> • 143, **144**, 149
- DENSE\_RANK • 63, 66, 147, 207, 208, 550, 1236
- depend immediately • 377
- dependent • 88, 352, 354, 561, 562, 565, 616, 620, 635, 636, 664, 713, 714, 738, 787, 800, 1223, 1286
- dependent on • 51, 71, 88, 93, 121, 203, 205, 358, 371, 561, 562, 564, 565, 566, 619, 620, 621, 624, 641, 663, 664, 712, 713, 714, 734, 735, 736, 738, 740, 743, 747, 788, 1223, 1247
- dependent privilege descriptor • 121
- dependent privilege descriptors still exist* • 800, 1147
- <dependent variable expression> • 65, 66, 205, 547, **548**, 550, 553, 554, 557, 558, 1174, 1175
- DEPTH • 146, 393, 394
- DEREF • 147, 244, 246
- <dereference operation> • 241, **243**, 372, 610, 622, 635, 792, 793, 795, 796, 798, 1190
- <dereference operator> • **241**, 243, 244
- DERIVED • 146, 574, 577, 628, 631
- <derived column> • 198, 277, 337, 344, 350, **366**, 369, 370, 371, 373, 375, 385, 386, 432, 860, 1020, 1041, 1174, 1214, 1219, 1249
- <derived column list> • 325, **326**, 328, 330, 331, 1261
- <derived representation> • 17, 41, 42, 47, 577, 581, 628, 629, **675**, 678, 685, 876, 884
- <derived table> • 61, 76, 325, **326**, 330, 331, 333, 334, 372, 379, 384, 1177
- DESC • 61, 62, 146, 361, 362, 363, 556, 559
- descending sequence generator • 82, 505
- DESCRIBE • 147, 1017
- describe • 15, 86, 121, 1018, 1076, 1130
- describe input • 86
- <describe input statement> • 86, 109, 111, **1017**, 1018, 1019, 1022, 1024, 1162, 1250, 1274
- describe output • 86
- <describe output statement> • 86, 109, 111, **1017**, 1018, 1019, 1022, 1028, 1161, 1250
- <describe statement> • 833, 998, 1002, **1017**, 1130, 1241
- <described object> • **1017**
- DESCRIPTOR • 146, 808, 993, 995, 996, 999, 1017, 1027, 1120
- descriptor • 6
- <descriptor item name> • **996**, 997, 998, 999, 1000, 1215
- <descriptor name> • **163**, 168, 171, 808, 993, 995, 996, 997, 998, 999, 1000, 1017, 1018, 1023, 1024, 1027, 1028, 1120, 1160, 1162, 1163, 1241, 1271
- determinant • 72
- DETERMINISTIC • 66, 147, 529, 677, 680, 686, 687, 688, 698, 717, 720, 725, 732, 734, 751, 1238
- deterministic • 40, 42, 66, 67, 75, 77, 78, 91, 92, 297, 301, 585, 601, 661, 685, 732, 736, 746, 750, 759, 1233
- <deterministic characteristic> • 676, 680, 698, 716, **717**, 720
- DIAGNOSTICS • 146, 945, 949, 950, 1125



- diagnostics area stack • 97, 131, 137, 524, 941, 1142, 1246
- diagnostics exception* • 1142, 1147
- <diagnostics size> • **949**, 950, 969, 970, 1167
- <digit> • **139**, 143, 149, 153, 157, 159, 160, 161, 176, 221, 223, 424, 893, 1067, 1143, 1176, 1243
- <direct implementation-defined statement> • **1119**, 1120, 1121, 1242
- <direct invocation> • **235**
- direct result of executing an SQL-statement • 11
- <direct select statement: multiple rows> • 107, 111, 114, 116, 1119, 1120, **1123**, 1131
- <direct SQL data statement> • **1119**, 1122
- <direct SQL statement> • 87, 91, 95, 123, 128, 129, 130, 133, 137, 164, 165, 166, 167, 487, 518, 519, 733, 834, 835, 974, 975, 977, 978, **1119**, 1120, 1121, 1122, 1160, 1227, 1251
- direct subtable • 58, 578, 623, 627
- direct subtype • 16, 41, 43, 58, 578, 628, 678
- direct subview • 627, 636
- direct supercolumn • 578
- direct supertable • 58, 578
- direct superview • 627
- directly based on • 16
- directly dependent • 787, 790
- <directly executable statement> • **1119**
- DISCONNECT • 147, 966
- disconnect error* • 967, 1151
- <disconnect object> • **966**
- <disconnect statement> • 108, 128, 832, **966**, 967, 1131, 1183, 1250
- DISPATCH • 146, 688, 716, 751
- <dispatch clause> • **716**, 721
- DISPLAY • 1088
- DISTINCT • 23, 29, 50, 64, 65, 147, 208, 209, 230, 252, 308, 309, 310, 311, 353, 367, 370, 373, 374, 375, 376, 380, 381, 382, 385, 386, 388, 389, 391, 408, **441**, 443, 481, 483, **547**, 549, 550, 552, 557, 558, 675, 1157, 1169, 1174, 1177, 1184, 1201, 1208, 1210, 1211, 1218, 1219, 1248
- distinct • **8**
- <distinct predicate> • 251, 401, 402, **441**, 443, 478, 479, 1009, 1210, 1211
- <distinct predicate part 2> • 212, **441**, 443, 1211
- distinct type • **8**, 13, 41, 45, 60, 477, 586, 605, 676, 684, 686, 750, 919, 926
- distinct types • 40, 1280
- DISTINT • 1157
- division by zero* • 255, 262, 1145
- does not possibly contain SQL • 40, 91, 525, 530, 732, 734, 736
- DOMAIN • 146, 571, 638, 641, 646, 647, 748, 754, 772, 779, 801
- <domain constraint> • 51, 192, 229, **638**, 639, 644
- <domain definition> • 105, 167, 177, 567, 569, 589, 601, **638**, 639, 640, 789, 831, 1131, 1168, 1179, 1249
- <domain name> • 51, 71, **161**, 167, 169, 171, 202, 215, 217, 229, 372, 571, 584, 585, 586, 587, 588, 592, 638, 639, 641, 642, 643, 644, 645, 646, 748, 754, 772, 779, 780, 789, 801, 1007, 1168
- dormant • 128, 129
- dormant SQL-connection • 128, 962
- dormant SQL-session • 129, 962, 964
- DOUBLE • 14, 15, 30, 147, 174, 178, 179, 183, 469, 474, 840, 842, 843, 844, 846, 847, 984, 986, 1076, 1085, 1094, 1099, 1230, 1231
- <double colon> • **145**, 237
- <double period> • **145**, 1071, 1105
- <double quote> • 20, 139, **140**, 144, 148, 149, 150, 282, 286
- <doublequote symbol> • **144**, 150, 286
- DROP • 147, 570, 571, 608, 610, 611, 616, 617, 619, 620, 621, 622, 623, 624, 635, 636, 637, 643, 645, 646, 650, 654, 659, 663, 664, 671, 695, 708, 712, 713, 714, 743, 744, 747, 748, 753, 754, 761, 762, 763, 764, 765, 769, 785, 801, 802, 898, 1165
- <drop assertion statement> • 105, 571, 611, 620, 624, 637, **663**, 664, 665, 748, 754, 801, 832, 1131, 1176
- <drop attribute definition> • 692, **695**, 696
- <drop behavior> • **570**, 610, 616, 619, 622, 624, 635, 637, 646, 654, 663, 712, 714, 743, 744, 747, 753, 761, 763, 769, 786, 802, 1165, 1267
- <drop character set statement> • 105, 166, 571, **650**, 651, 832, 1131, 1175
- <drop collation statement> • 105, 166, 571, **654**, 655, 802, 832, 1131, 1179
- <drop column default clause> • 606, **608**, 1170, 1171
- <drop column definition> • 603, **616**, 617, 1165
- <drop column scope clause> • 606, **610**, 611, 1171, 1191
- <drop data type statement> • 88, 106, 571, **712**, 714, 802, 832, 1131, 1165, 1273
- <drop domain constraint definition> • 641, **645**, 1176, 1181
- <drop domain default clause> • 641, **643**, 1180
- <drop domain statement> • 105, 571, **646**, 647, 748, 754, 801, 832, 1131, 1168
- <drop method specification> • 692, **708**, 711
- <drop role statement> • 106, 121, **785**, 832, 1131, 1216
- <drop routine statement> • 88, 106, 571, 610, 620, 624, 637, 664, 713, **743**, 744, 748, 753, 762, 764, 765, 802, 832, 1131, 1165, 1188, 1273
- <drop schema statement> • 105, **570**, 572, 831, 1131, 1170

<drop sequence generator statement> • 106, 571, **769**, 832, 1131, 1212  
 <drop table constraint definition> • 603, **619**, 621, 1171  
 <drop table statement> • 105, 570, **622**, 623, 624, 748, 801, 831, 898, 1131, 1165, 1267  
 <drop transform element list> • 758, **761**, 762  
 <drop transform statement> • 106, 713, 744, **763**, 765, 832, 1131, 1199  
 <drop transliteration statement> • 106, 166, 571, **659**, 660, 832, 1131, 1180  
 <drop trigger statement> • 106, 571, 617, 621, 624, 637, 664, **671**, 673, 748, 754, 801, 832, 1131, 1212  
 <drop user-defined cast statement> • 106, 713, 714, 744, **747**, 748, 832, 1131, 1195, 1196  
 <drop user-defined ordering statement> • 106, 744, **753**, 754, 832, 1131, 1200  
 <drop view statement> • 105, 570, 611, 620, 624, **635**, 636, 637, 664, 748, 754, 801, 831, 1131, 1165, 1267  
 DT-EC-ordered • 17  
 DT-FC-ordered • 17  
 DT-NC-ordered • 17  
 duplicates • **8**, 50, 64, 65, 389, 390, 443, 448, 553, 1248  
 dyadic • **8**  
 DYNAMIC • 147, 716, 720, 805, 807, 808, 1056, 1163  
 <dynamic close statement> • 101, 107, 111, 113, 116, 117, 833, **1043**, 1131, 1161  
 dynamic cursor • 10, 86, 99, 102, 170, 902, 1028, 1039, 1041, 1043, 1044, 1046, 1048, 1052, 1241  
 <dynamic cursor name> • **163**, 170, 171, 1039, 1041, 1043, 1044, 1046, 1160, 1242  
 <dynamic declare cursor> • 10, 86, 99, 100, 102, 107, 112, 114, 116, 117, 165, 805, 807, 849, 853, 854, **1035**, 1055, 1058, 1061, 1065, 1161  
 <dynamic delete statement: positioned> • 102, 103, 107, 111, 114, 116, 117, 833, 870, 872, 882, 890, 909, 911, 938, **1044**, 1045, 1131, 1161  
 <dynamic fetch statement> • 86, 94, 102, 107, 111, 113, 116, 117, 833, 1002, 1027, 1028, **1041**, 1131, 1161, 1241  
 <dynamic open statement> • 86, 101, 107, 111, 113, 115, 117, 833, 853, 854, 1002, 1023, **1039**, 1040, 1131, 1161, 1241, 1255  
 <dynamic parameter specification> • 85, 90, 190, **191**, 192, 193, 194, 294, 371, 466, 517, 520, 538, 574, 584, 601, 625, 661, 668, 723, 899, 901, 902, 1004, 1011, 1012, 1017, 1018, 1019, 1020, 1021, 1022, 1023, 1024, 1027, 1028, 1033, 1034, 1039, 1119, 1160, 1162, 1241, 1250  
 <dynamic select statement> • 86, 107, 112, 114, 115, 116, 117, 834, 1003, **1004**, 1018, 1019, 1023, 1027, 1028, 1032, 1034, 1130, 1131

<dynamic single row select statement> • 86, 107, 111, 114, 115, 116, 117, 230, 834, 1003, 1004, 1018, 1019, 1027, 1028, 1032, 1033, 1034, **1042**, 1130, 1131, 1161  
 dynamic SQL argument list • 522  
 dynamic SQL error • 993, 997, 998, 1000, 1001, 1002, 1023, 1024, 1025, 1026, 1027, 1028, 1029, 1030, 1032, 1033, 1037, 1039, 1147, 1241, 1255  
 <dynamic update statement: positioned> • 102, 103, 107, 111, 114, 116, 117, 833, 870, 872, 882, 890, 909, 911, 938, **1046**, 1047, 1132, 1161, 1162  
 DYNAMIC\_FUNCTION • 146, 991, 996, 999, 1018, 1125, 1129, 1139  
 DYNAMIC\_FUNCTION\_CODE • 146, 991, 996, 999, 1018, 1125, 1129

## — E —

E • 153  
 EACH • 147, 666, 668, 669, 670, 1212  
 effective binary precision • 469  
 effective returns data type • 519, 521, 522, 535  
 effective SQL parameter list • 90, 93, 528, 529, 726, 728, 729, 730, 731, 737, 741  
 effectively updatable • 626  
 ELEMENT • 147, 249  
 element order • 49  
 element type • **8**, 16, 47, 48, 49, 179, 180, 183, 196, 217, 251, 252, 304, 306, 308, 309, 312, 313, 327, 442, 444, 446, 448, 454, 458, 470, 484, 492, 498, 528, 536, 551, 553, 693, 695, 712, 721, 726, 727, 859, 860, 864, 894, 895, 983, 986, 1188  
 elements • xix, xx, xx, xx, xx, 3, 7, 47, 48, 49, 52, 61, 85, 122, 136, 139, 219, 262, 307, 313, 324, 378, 387, 404, 406, 442, 445, 447, 457, 462, 470, 477, 553, 557, 605, 704, 705, 709, 710, 736, 862, 864, 1002, 1227, 1246, 1247, 1249, 1250, 1257  
 ELSE • 147, 208, 211, 212, 213, 308, 309, 311, 327, 395, 1140  
 <else clause> • **211**, 213, 214  
 <embedded authorization clause> • 1055, **1056**, 1060  
 <embedded authorization declaration> • 85, **1055**, 1058, 1060, 1065, 1163, 1241, 1251  
 <embedded authorization identifier> • **1056**  
 <embedded character set declaration> • **1056**, 1058, 1059, 1064, 1065, 1175, 1242  
 <embedded collation specification> • 1055, **1056**, 1058, 1060  
 <embedded exception declaration> • 109, 1055, 1061, 1065, **1067**, 1068, 1069, 1070, 1176, 1271  
 <embedded path specification> • 1055, **1056**, 1058, 1060, 1066, 1192, 1242  
 <embedded SQL Ada program> • 1055, 1057, 1059, 1065, 1068, **1071**, 1072, 1073, 1075, 1076, 1159, 1242

- <embedded SQL begin declare> • **1056**, 1057, 1058, 1064, 1106
- <embedded SQL C program> • 1055, 1057, 1059, 1065, 1068, **1078**, 1080, 1081, 1084, 1085, 1159, 1242
- <embedded SQL COBOL program> • 1055, 1057, 1060, 1065, 1068, **1087**, 1088, 1089, 1091, 1092, 1159, 1242
- <embedded SQL declare section> • 84, **1056**, 1057, 1065, 1072, 1080, 1088, 1095, 1101, 1106, 1112
- <embedded SQL end declare> • **1056**, 1057, 1058, 1064, 1106
- <embedded SQL Fortran program> • 1055, 1057, 1060, 1065, 1068, **1094**, 1095, 1098, 1099, 1159, 1242
- <embedded SQL host program> • 84, 85, **1055**, 1057, 1058, 1059, 1065, 1068, 1069, 1242, 1251, 1271, 1287
- <embedded SQL MUMPS declare> • **1056**, 1057, 1064
- <embedded SQL MUMPS program> • 1055, 1057, 1060, 1068, **1100**, 1101, 1103, 1104, 1159, 1206
- <embedded SQL Pascal program> • 1055, 1057, 1060, 1065, 1068, **1105**, 1106, 1109, 1160, 1242
- <embedded SQL PL/I program> • 1055, 1057, 1060, 1065, 1068, 1069, **1111**, 1112, 1115, 1116, 1160, 1242
- <embedded SQL statement> • 84, 85, 192, **1055**, 1057, 1058, 1059, 1065, 1072, 1080, 1088, 1089, 1095, 1101, 1106, 1112
- <embedded transform group specification> • 1055, **1056**, 1058, 1060, 1061, 1063, 1066, 1199
- <embedded variable name> • 190, 191, 192, 193, 194, 668, 723, 899, **1056**, 1060, 1061, 1065, 1251, 1271
- <embedded variable specification> • 190, **191**, 193, 194, 520, 535, 574, 584, 601, 625, 661, 858, 860, 863, 865, 901, 902, 1023, 1027, 1119
- empty • 45, 54, 55, 56, 69, 70, 71, 72, 73, 77, 96, 97, 101, 134, 135, 136, 216, 218, 239, 241, 244, 340, 341, 387, 388, 425, 430, 431, 469, 494, 518, 522, 526, 536, 552, 557, 579, 611, 624, 630, 636, 669, 685, 712, 741, 756, 834, 863, 877, 883, 884, 890, 905, 906, 907, 912, 913, 917, 920, 921, 923, 925, 926, 927, 928, 930, 1020, 1037, 1042, 1123
- <empty grouping set> • **345**, 346, 347, 348, 349, 350, 351, 353, 1217
- <empty specification> • **196**, 197, 215, 218, 316, 589, 590, 591, 875, 880, 894, 1193, 1200
- emptying • 97
- enabled • 122, 570, 580, 603, 623, 629, 636, 639, 641, 646, 648, 650, 652, 654, 657, 659, 661, 664, 668, 671, 684, 692, 713, 732, 741, 744, 746, 748, 751, 753, 756, 758, 760, 764, 767, 768, 769, 785, 871, 1137, 1138
- END • 147, 208, 211, 212, 308, 309, 311, 327, 395, 666, 1056, 1064, 1074, 1107, 1140
- <end field> • 179, 296, 464, **509**, 510, 511, 512
- END-EXEC • 147
- enduring characteristics • 131
- enduring transaction characteristics • 131, 945, 956, 958, 963
- equal • 8
- equality AND-component • 73
- equality operation • 337, 413, 436, 444, 478, 479, 480, 483, 753, 754, 1177, 1202
- EQUALS • 17, 41, 44, 146, 687, 749, 750, 751, 753
- <equals operator> • 8, 23, 29, 73, 140, **141**, 403, 410, 431, 478, 483, 485, 892, 996, 999, 1071, 1080, 1125, 1126, 1268
- <equals ordering form> • **749**, 750
- error in assignment* • 997, 1000, 1145
- ESCAPE • 147, 273, 283, 414, 415, 416, 417, 420, 421, 422, 423, 1005
- <escape character> • 273, 276, 282, **414**, 416, 419, 420, 421, 422, 478, 479, 1009, 1169, 1261
- escape character conflict* • 422, 1145
- <escape octet> • **414**, 415, 417, 418, 478, 1009
- <escaped character> • 420, **421**, 422
- EVERY • 65, 147, 547, 549, 552, 557, 588, 594, 595, 599, 1204
- <exact numeric literal> • 31, **153**, 157, 159, 221, 222, 228, 590, 1229, 1258
- <exact numeric type> • 30, 31, **174**, 178, 182, 820, 1021, 1076, 1230, 1258
- exact numeric types • 15
- EXCEPT • 23, 29, 50, 79, 80, 147, 251, 252, 308, 309, 310, 376, 379, 380, 381, 382, 383, 384, 385, 388, 389, 390, 391, 408, 481, 483, 924, 931, 1169, 1201, 1219, 1267, 1272
- exception • 95
- exception condition • 18, 28, 83, 86, 95, 97, 98, 102, 123, 126, 128, 136, 180, 193, 194, 220, 221, 222, 223, 224, 225, 226, 227, 228, 229, 234, 248, 255, 262, 263, 264, 265, 270, 271, 282, 290, 297, 333, 361, 398, 399, 416, 417, 422, 455, 456, 457, 459, 460, 461, 462, 490, 493, 496, 504, 522, 525, 526, 527, 530, 531, 532, 534, 535, 536, 551, 552, 555, 557, 691, 776, 787, 800, 804, 813, 820, 826, 835, 836, 837, 859, 860, 863, 864, 865, 869, 871, 876, 881, 882, 889, 895, 896, 901, 903, 905, 907, 908, 909, 910, 911, 924, 931, 934, 935, 936, 937, 941, 945, 947, 948, 953, 954, 955, 957, 958, 961, 962, 964, 966, 971, 972, 973, 974, 975, 977, 978, 979, 980, 993, 995, 997, 998, 1000, 1001, 1002, 1011, 1012, 1013, 1015, 1018, 1023, 1024, 1025, 1026, 1027, 1028, 1029, 1030, 1032, 1033, 1034, 1036, 1037, 1039, 1041, 1042, 1044, 1046, 1047, 1120, 1121, 1122, 1136, 1142, 1143, 1152, 1231, 1241, 1249, 1250
- exception data item • 490, 491, 528, 529, 531
- EXCLUDE • 146, 357, 365
- EXCLUDING • 146, 574, 575
- <exclusive user-defined type specification> • **449**
- exclusively specified • 449

EXEC • 147, 1055, 1057  
 executable routine • 517  
 EXECUTE • 120, 121, 147, 218, 240, 517, 631, 632, 657, 732, 733, 744, 771, 772, 773, 774, 779, 780, 792, 793, 794, 796, 797, 1032, 1034, 1065  
 <execute immediate statement> • 11, 83, 86, 95, 109, 111, 112, 129, 130, 164, 165, 166, 167, 487, 518, 519, 833, 974, 975, 977, 978, **1034**, 1132, 1161, 1229, 1251  
 execute privilege descriptor • 121  
 <execute statement> • 11, 86, 109, 111, 112, 833, 978, 1002, 1023, 1027, **1032**, 1033, 1132, 1161, 1162, 1241  
 executed immediately • 85, 86, 116  
 executing statement • 128, 526, 813, 834, 837, 941, 1033  
 <existing collation name> • **652**, 797, 1237  
 <existing transliteration name> • 284, **656**, 657, 1237  
 <existing window name> • **356**, 358, 359, 360, 365, 1221  
 EXISTS • 147, 432, 596, 924, 931, 1218  
 <exists predicate> • 23, 29, 367, 401, 402, **432**, 616, 1218, 1262  
 EXP • 147, 258, 263  
 expandable • 380, 393  
 <explicit row value constructor> • **315**, 316, 317, 318, 319, 1178, 1208  
 <explicit table> • 54, **376**, 380, 381, 384, 391, 565, 1178, 1248  
 <exponent> • 31, **153**, 159  
 <exponential function> • 32, 256, **258**, 260, 263, 266, 1222, 1232  
 exposed • 199, 200, 202, 329, 330, 333, 367, 368, 372, 855, 866, 868, 869, 879, 881, 883, 886, 888, 889, 1052  
 <extended cursor name> • **163**, 168, 170, 171, 1017, 1018, 1036, 1037, 1038, 1050, 1162  
 <extended descriptor name> • **164**, 168, 170, 993  
 extended dynamic cursor • 99, 170, 537, 1018, 1050  
 extended names • 87  
 extended order by clause • 386, 390  
 extended sort key columns • 386, 390  
 extended statement cursor • 86, 99, 100, 101, 102, 901, 1037  
 <extended statement name> • **163**, 168, 170, 171, 853, 1012, 1036, 1037, 1162  
 EXTERNAL • 147, 717, 725, 737, 1237  
 <external body reference> • 89, **717**, 720  
 external routine • **8**, 45, 89, 90, 91, 92, 93, 129, 453, 454, 523, 525, 526, 527, 537, 713, 720, 721, 722, 723, 725, 731, 732, 733, 734, 736, 741, 761, 762, 763, 764, 1135, 1136, 1235  
 external routine exception • 525, 530, 531, 1138, 1147  
 external routine invocation exception • 529, 1138, 1148  
 <external routine name> • 89, **162**, 169, 717, 725, 734, 736, 738, 740, 741, 742, 1222, 1223

external routine SQL-path • 91  
 <external security clause> • **717**, 725, 737, 1215  
 external sequence generator • 82  
 <externally-invoked procedure> • 1, 11, 55, 83, 84, 85, 87, 93, 95, 98, 100, 104, 113, 114, 118, 119, 123, 128, 129, 137, 568, 805, 806, 807, 808, **811**, 812, 813, 814, 823, 824, 834, 835, 836, 849, 902, 961, 962, 965, 966, 967, 1013, 1060, 1061, 1064, 1065, 1121, 1227, 1228, 1240, 1251, 1283, 1287  
 EXTRACT • 147, 257, 261  
 <extract expression> • 32, 39, 256, **257**, 259, 261, 266, 1166, 1173, 1232  
 <extract field> • **257**, 259, 261  
 <extract source> • **257**, 259, 260, 261

## — F —

Feature F032, "CASCADE drop behavior" • 624, 637, 714, 744, 1165  
 Feature F033, "ALTER TABLE statement: DROP COLUMN clause" • 617, 1165  
 Feature F034, "Extended REVOKE statement" • 802, 803, 1165, 1166  
 Feature F052, "Intervals and datetime arithmetic" • 160, 184, 266, 290, 297, 298, 512, 1166  
 Feature F053, "OVERLAPS predicate" • 440, 1166, 1167  
 Feature F111, "Isolation levels other than SERIALIZABLE" • 950, 1167  
 Feature F121, "Basic diagnostics management" • 950, 1140, 1167  
 Feature F122, "Enhanced diagnostics management" • 1141, 1167  
 Feature F123, "All diagnostics" • 1141, 1167  
 Feature F131, "Grouped operations" • 1258, 1260, 1264, 1270  
 Feature F171, "Multiple schemas per user" • 569, 1167  
 Feature F191, "Referential delete actions" • 600, 1167  
 Feature F200, "TRUNCATE TABLE statement" • 872, 1167, 1168  
 Feature F222, "INSERT statement: DEFAULT VALUES clause" • 877, 1168  
 Feature F251, "Domain support" • 171, 194, 639, 647, 1168  
 Feature F262, "Extended CASE expression" • 214, 1168  
 Feature F263, "Comma-separated predicates in simple CASE expression" • 214, 1168  
 Feature F271, "Compound character literals" • 160, 1168, 1169  
 Feature F281, "LIKE enhancements" • 419, 1169  
 Feature F291, "UNIQUE predicate" • 433, 1169  
 Feature F301, "CORRESPONDING in query expressions" • 391, 1169

- Feature F302, "INTERSECT table operator" • 391, 1169
- Feature F304, "EXCEPT ALL table operator" • 391, 1169
- Feature F312, "MERGE statement" • 885, 1169
- Feature F313, "Enhanced MERGE statement" • 885, 1169, 1170
- Feature F321, "User authorization" • 194, 592, 971, 1170
- Feature F361, "Subprogram support" • 1066, 1170
- Feature F381, "Extended schema manipulation" • 572, 606, 607, 608, 609, 611, 618, 621, 742, 1170, 1171
- Feature F382, "Alter column data type" • 614, 1171
- Feature F391, "Long identifiers" • 151, 1171
- Feature F392, "Unicode escapes in identifiers" • 151, 1171
- Feature F393, "Unicode escapes in literals" • 160, 1171, 1172
- Feature F394, "Optional normal form specification" • 287, 435, 1172
- Feature F401, "Extended joined table" • 343, 1172
- Feature F402, "Named column joins for LOBs, arrays, and multisets" • 343, 1172
- Feature F403, "Partitioned join tables" • 343, 1172
- Feature F411, "Time zone specification" • 160, 184, 266, 290, 292, 973, 1172, 1173
- Feature F421, "National character" • 160, 184, 229, 266, 419, 1173, 1207
- Feature F431, "Read-only scrollable cursors" • 852, 860, 1173, 1174
- Feature F441, "Extended set function support" • 344, 557, 1174
- Feature F442, "Mixed column references in set functions" • 558, 1174, 1175
- Feature F451, "Character set definition" • 649, 651, 1175
- Feature F461, "Named character sets" • 171, 540, 569, 810, 977, 1065, 1175
- Feature F491, "Constraint management" • 171, 546, 645, 1070, 1176
- Feature F521, "Assertions" • 662, 665, 1176
- Feature F531, "Temporary tables" • 582, 898, 1176
- Feature F555, "Enhanced seconds precision" • 160, 184, 292, 1176, 1177
- Feature F561, "Full value expressions" • 413, 557, 1177
- Feature F571, "Truth value tests" • 303, 1177
- Feature F591, "Derived tables" • 334, 1177
- Feature F611, "Indicator data types" • 194, 1177, 1231
- Feature F641, "Row and table constructors" • 317, 321, 1177, 1178
- Feature F651, "Catalog name qualifiers" • 171, 974, 1178
- Feature F661, "Simple tables" • 391, 1178
- Feature F671, "Subqueries in CHECK constraints" • 602, 1178
- Feature F672, "Retrospective check constraints" • 602, 662, 1178, 1179
- Feature F690, "Collation support" • 171, 544, 653, 655, 1179
- Feature F692, "Extended collation support" • 588, 640, 691, 1179
- Feature F693, "SQL-session and client module collations" • 195, 808, 981, 1179, 1180
- Feature F695, "Translation support" • 171, 287, 658, 660, 1180
- Feature F701, "Referential update actions" • 600, 1180
- Feature F711, "ALTER domain" • 641, 642, 643, 644, 645, 1180, 1181
- Feature F721, "Deferrable constraints" • 546, 952, 1181
- Feature F731, "INSERT column privileges" • 781, 1181
- Feature F741, "Referential MATCH types" • 438, 600, 1181
- Feature F751, "View CHECK enhancements" • 633, 1181
- Feature F761, "Session management" • 970, 974, 976, 977, 1181, 1182
- Feature F762, "CURRENT\_CATALOG" • 194, 592, 1182
- Feature F763, "CURRENT\_SCHEMA" • 194, 592, 1182
- Feature F771, "Connection management" • 171, 963, 965, 967, 1182, 1183
- Feature F781, "Self-referencing operations" • 870, 877, 884, 885, 891, 896, 1183
- Feature F791, "Insensitive cursors" • 852, 1183, 1213
- Feature F801, "Full set function" • 374, 1183, 1184
- Feature F821, "Local table references" • 171, 203, 1184
- Feature F831, "Full cursor update" • 850, 855, 887, 1184
- Feature F841, "LIKE\_REGEX predicate" • 427, 1184
- Feature F842, "OCCURENCES\_REGEX function" • 267, 1184
- Feature F843, "POSITION\_REGEX function" • 267, 1184
- Feature F844, "SUBSTRING\_REGEX" • 287, 1184, 1185
- Feature F845, "TRANSLATE\_REGEX" • 287, 1185
- Feature F846, "Octet support in regular expression operators" • 267, 287, 1185
- Feature F847, "Nonconstant regular expression" • 267, 287, 427, 1185
- Feature F850, "Top-level <order by clause> in <query expression>" • 391, 1185
- Feature F851, "<order by clause> in subqueries" • 391, 1185, 1186
- Feature F852, "Top-level <order by clause> in views" • 634, 1186
- Feature F855, "Nested <order by clause> in <query expression>" • 391, 1186
- Feature F856, "Nested <fetch first clause> in <query expression>" • 391, 1186
- Feature F857, "Top-level <fetch first clause> in <query expression>" • 392, 1186
- Feature F858, "<fetch first clause> in subqueries" • 392, 1186

Feature F859, "Top-level <fetch first clause> in views" • 634, 1186

<factor> • **254**, 293, 294, 295

FALSE • 147, 155, 160, 222, 224, 299, 301, 826, 828

*feature not supported* • 947, 961, 964, 1148

FETCH • 147, 377, 857, 1041

fetch call • 490

<fetch first clause> • 54, 376, **377**, 380, 381, 384, 386, 391, 392, 634, 1155, 1186, 1280

<fetch orientation> • **857**, 858, 860, 1041, 1174

<fetch statement> • 94, 102, 106, 110, 113, 116, 117, 807, 832, 836, **857**, 860, 909, 911, 1132, 1174, 1249, 1264, 1265

<fetch target list> • **857**, 858, 859, 1249

field • 52

<field definition> • 175, 183, **186**, 187, 1207

<field name> • 46, **163**, 169, 171, 186, 232, 316, 331, 336, 337, 338, 373, 464, 465, 582, 603, 893, 1207

<field reference> • 188, 189, 200, **232**, 372, 668, 1208

fields • xix, 13, 16, 34, 36, 37, 38, 39, 46, 52, 61, 86, 180, 181, 182, 183, 290, 293, 294, 323, 336, 337, 339, 342, 343, 403, 442, 464, 465, 477, 511, 512, 590, 628, 911, 925, 1001, 1004, 1019, 1233, 1247, 1265

FILTER • 147, 547

<filter clause> • 64, 205, **547**, 550, 551, 558, 1221

FINAL • 146, 675, 676, 677, 689, 1189

<finality> • **674**, 676, 677, 689, 1189, 1190

FIRST • 61, 62, 146, 361, 362, 363, 377, 393, 559, 560, 857, 905, 906

fixed-length • **8**

FLAG • 146, 256, 273, 274, 426

FLOAT • 14, 15, 30, 147, 174, 178, 179, 183, 469, 474, 840, 841, 843, 844, 845, 847, 984, 986, 1001, 1112, 1116, 1230, 1231

FLOOR • 147, 258

<floor function> • 32, 256, **258**, 260, 263, 267, 1222, 1232

<fold> • 22, **273**, 275, 276, 281, 283, 284, 1259, 1283

FOLLOWING • 146, 208, 357, 358, 361, 362, 363, 364

FOR • 87, 101, 147, 191, 230, 261, 273, 274, 275, 278, 280, 538, 541, 652, 656, 666, 668, 669, 670, 672, 673, 686, 687, 713, 716, 717, 744, 749, 753, 755, 758, 763, 786, 790, 791, 800, 802, 803, 805, 807, 808, 849, 850, 853, 854, 855, 904, 979, 980, 1005, 1013, 1035, 1036, 1056, 1163, 1165, 1184, 1193, 1209, 1212

FOREIGN • 147, 586, 597

FORTRAN • 146, 489, 513, 529, 731, 739, 741, 809, 823, 825, 826, 827, 828, 829, 1059, 1163, 1164

<Fortran array locator variable> • 1094, **1095**, 1098, 1099, 1197

<Fortran BINARY variable> • **1094**, 1096, 1099, 1203

<Fortran BLOB locator variable> • 1094, **1095**, 1097, 1099, 1206

<Fortran BLOB variable> • 1094, **1095**, 1096, 1099, 1205

<Fortran CLOB locator variable> • 1094, **1095**, 1097, 1099, 1206

<Fortran CLOB variable> • **1094**, 1096, 1099, 1205

<Fortran derived type specification> • **1094**

<Fortran host identifier> • 1057, **1094**, 1095, 1096, 1098

<Fortran multiset locator variable> • 1094, **1095**, 1098, 1099, 1198

<Fortran REF variable> • 1094, **1095**, 1098, 1099, 1190

<Fortran type specification> • **1094**, 1095, 1096

<Fortran user-defined type locator variable> • 1094, **1095**, 1097, 1099, 1196

<Fortran user-defined type variable> • 1094, **1095**, 1097, 1099, 1199

<Fortran VARBINARY variable> • **1094**, 1097, 1099, 1203

<Fortran variable definition> • 1056, **1094**, 1095, 1098

FOUND • 146, 1067, 1068, 1070

FREE • 147, 899

<free locator statement> • 97, 107, 110, 111, 114, 115, 117, 725, 832, **899**, 1003, 1011, 1132, 1219, 1238, 1283

FROM • 59, 75, 76, 78, 147, 219, 224, 225, 226, 227, 228, 243, 247, 249, 256, 257, 261, 273, 274, 275, 278, 279, 280, 308, 309, 311, 312, 323, 328, 329, 336, 337, 341, 342, 343, 352, 370, 380, 382, 389, 394, 395, 396, 441, 555, 556, 582, 588, 594, 595, 596, 599, 617, 623, 626, 633, 647, 651, 652, 655, 656, 659, 671, 672, 673, 675, 687, 714, 716, 744, 755, 761, 769, 785, 786, 857, 866, 868, 883, 920, 924, 931, 955, 961, 962, 971, 972, 974, 975, 977, 978, 979, 980, 993, 1003, 1005, 1012, 1034, 1036, 1041, 1044, 1048, 1134, 1135

<from clause> • v, 59, 77, 202, 209, 322, **323**, 324, 329, 344, 346, 352, 359, 369, 372, 373, 379, 394, 563, 628, 915, 921, 928, 1270

<from constructor> • 230, **873**, 875

<from default> • **873**, 877, 1168

<from sql> • **755**, 756, 757, 759, 760, 761

<from sql function> • **755**, 756, 759, 760

<from subquery> • 230, **873**, 874, 875, 877, 1183

from-sql function • 45, 93, 500, 501, 524, 726, 729, 732, 734, 755, 759, 761, 762, 763, 764, 1030

from-sql function associated with the i-th SQL parameter • 726, 729, 730, 731

FULL • 17, 41, 44, 68, 69, 74, 75, 147, 336, 338, 339, 341, 343, 379, 380, 436, 437, 485, 597, 599, 687, 749, 750, 751, 752, 753, 933, 934, 935, 936, 1172, 1200

<full ordering form> • **749**

fully qualified • **6**

FUNCTION • 42, 147, 541, 542, 543, 686, 687, 688, 715, 751

*function executed no return statement* • 527, 1150

<function specification> • 89, **715**, 720, 724, 1273  
 functional dependency • v, 71, 72, 74, 75, 76, 77, 78, 80, 371  
 FUSION • 147, 547, 549, 550, 553, 558, 1202

## — G —

G • 143, 146, 176  
 GENERAL • 90, 146, 529, 532, 534, 688, 717, 730, 734, 736, 741, 742, 1188  
 <general literal> • **152**  
 <general set function> • 23, 29, 205, 483, 484, **547**, 548, 549, 551, 557, 558, 1174, 1175, 1177  
 <general value specification> • 71, 129, 130, **190**, 193, 194, 588, 826, 827, 1023, 1160, 1168, 1170, 1182, 1192, 1231  
 <generalized expression> • **516**, 519, 520, 521, 522, 523, 538, 1187  
 <generalized invocation> • **235**, 236  
 generally underlying table • 57  
 GENERATED • 60, 146, 574, 575, 576, 577, 582, 584, 627, 628, 631, 675, 1191  
 generated column • 61, 599, 880  
 <generation clause> • **584**, 585, 587, 588, 1211  
 <generation expression> • 53, 61, 576, **584**, 585, 612, 616, 617, 920, 926  
 <generation option> • **574**, 575  
 <generation rule> • **584**  
 GET • 147, 648, 996, 1125  
 <get descriptor information> • **996**, 998  
 <get descriptor statement> • 86, 109, 111, 833, **996**, 997, 998, 1132, 1160, 1250  
 <get diagnostics statement> • 98, 109, 833, **1125**, 1127, 1129, 1136, 1140, 1141, 1152, 1167, 1251, 1255  
 <get header information> • **996**, 997, 998  
 <get item information> • **996**, 997, 998  
 GLOBAL • 55, 100, 147, 164, 573, 580, 1037, 1038, 1050  
 global • 54, 55, 56, 57, 87, 99, 129, 524, 525, 537, 570, 573, 580, 598, 601, 1050, 1228, 1235, 1245  
 <global or local> • **573**  
 GO • 146, 1067, 1068, 1069, 1070  
 <go to> • **1067**, 1068, 1069, 1070  
 GOTO • 146, 1067  
 <goto target> • **1067**, 1068  
 GRANT • 121, 147, 539, 582, 633, 648, 672, 673, 688, 772, 773, 774, 776, 777, 783, 786, 789, 790, 791, 800, 802, 1165  
 <grant privilege statement> • 105, 121, 771, **776**, 778, 1132, 1189, 1193, 1263  
 <grant role statement> • 106, 122, 771, **783**, 784, 1132, 1216

<grant statement> • 539, 568, 582, 633, 648, **771**, 772, 773, 774, 777, 780, 831, 1272  
 GRANTED • 146, 776, 777, 783, 786, 787  
 <grantee> • 776, 777, **779**, 780, 783, 784, 786, 787, 802, 803, 1166  
 <grantor> • 776, **779**, 780, 781, 782, 783, 786, 787, 804, 1217  
 <greater than operator> • 140, **141**, 403, 404, 560  
 <greater than or equals operator> • 144, **145**, 403, 404  
 GROUP • 147, 256, 273, 345, 346, 347, 350, 351, 352, 353, 354, 357, 365, 371, 548, 717, 979, 1218  
 group aggregate function • 63, 64  
 <group by clause> • v, 23, 29, 55, 78, 209, 322, 329, **345**, 346, 347, 349, 352, 353, 354, 359, 369, 371, 373, 481, 619, 628, 1217, 1218, 1247, 1261, 1270  
 <group name> • 717, 722, 726, 727, 729, 730, 731, 744, **755**, 756, 757, 758, 759, 761, 763, 806, 979, 1058, 1061, 1063  
 <group specification> • **717**, 722, 726, 727, 729, 730, 731, 806, 1058, 1061, 1063  
 group-invariant column reference • 203  
 grouped table • 55, 62, 64, 80, 322, 345, 352, 354, 357, 358, 365, 369, 371, 374, 386, 549, 1214  
 grouped view • 1270  
 grouped, windowed query • 369  
 GROUPING • 63, 78, 147, 205, 206, 345, 347, 348, 349, 350, 351, 353, 1157, 1217, 1218  
 grouping column • 206, 346, 352  
 <grouping column reference> • **345**, 346, 347, 349, 351  
 <grouping column reference list> • 78, **345**, 346, 351, 353, 663, 1217  
 <grouping element> • **345**, 349, 353, 1217  
 <grouping element list> • **345**, 350  
 <grouping operation> • 62, 63, **205**, 206, 308, 309, 311, 347, 351, 357, 367, 446, 479, 480, 481, 482, 483, 484, 550, 595, 598, 753, 754, 1188, 1202, 1217, 1218  
 <grouping set> • **345**, 347, 348, 349, 350  
 <grouping set list> • **345**, 353, 1217  
 <grouping sets specification> • **345**, 346, 347, 348, 349, 350, 353, 1217

## — H —

HAVING • 147, 354  
 <having clause> • v, 63, 78, 203, 205, 209, 322, 329, 336, 344, 346, **354**, 359, 369, 371, 373, 379, 628, 1261, 1270  
 <header item name> • **996**, 997, 998, 999, 1002  
*held cursor requires same isolation level* • 947, 1149  
 <hexit> • 144, 149, **153**, 155, 156, 157, 160, 1204  
 HIERARCHY • 120, 146, 246, 332, 582, 632, 633, 771, 773, 775, 776, 777, 778, 786, 790, 791, 792, 793, 795, 796, 799, 800, 803, 869, 880, 881, 888, 889, 1193

<high value> • **420**, 422, 423  
 HOLD • 94, 100, 102, 147, 851, 852, 854, 900, 1038, 1219  
 <hold locator statement> • 97, 107, 110, 111, 114, 115, 118, 725, 832, **900**, 1003, 1011, 1132, 1219, 1238  
 holdable • 97, 101, 102, 103, 129, 808, 900, 903, 909, 911, 947, 955, 1227  
 holdable locator • 97  
 <host identifier> • **1056**, 1059, 1065  
 <host label identifier> • **1067**, 1068, 1069, 1070  
 host language • xx, 1, 15, 34, 45, 83, 84, 85, 87, 90, 95, 98, 99, 513, 838, 1057, 1058, 1059, 1062, 1064, 1065, 1068, 1069, 1070, 1227, 1242, 1267, 1273, 1287  
 <host parameter data type> • 806, **811**, 813, 1060, 1061, 1196, 1197, 1198  
 <host parameter declaration> • 95, 806, 807, **811**, 812, 814, 822, 823, 834, 849, 1060, 1061, 1062, 1065, 1251  
 <host parameter declaration list> • 806, **811**  
 <host parameter name> • **162**, 169, 190, 191, 193, 277, 472, 661, 668, 723, 811, 812, 834, 849, 860, 899, 900, 1060, 1061, 1062, 1251  
 <host parameter specification> • 190, **191**, 192, 193, 277, 472, 520, 535, 574, 584, 601, 625, 858, 863, 865, 901, 902, 1023, 1027  
 <host PL/I label variable> • **1067**, 1068, 1069, 1070  
 <host variable definition> • **1056**, 1058, 1059, 1065, 1066, 1170  
 HOUR • 35, 129, 147, 158, 181, 261, 289, 464, 509, 510, 973, 1002, 1010, 1011  
 <hours value> • **154**, 158  
 <hypothetical set function> • 64, 207, **548**, 550, 555, 558, 1221, 1236  
 <hypothetical set function value expression list> • **548**, 550, 555, 558, 1174  
 HZ • 1149



identical • 5, 7, **8**, 10, 29, 42, 43, 48, 59, 66, 67, 91, 93, 191, 339, 476, 477, 529, 553, 598, 653, 709, 775, 777, 783, 909, 911, 924, 931, 937, 961, 1035, 1059, 1060, 1069, 1070, 1240  
 identified • 9, 17, 20, 22, 40, 45, 46, 47, 51, 54, 67, 71, 74, 79, 81, 82, 86, 87, 89, 90, 98, 100, 120, 121, 128, 136, 151, 156, 165, 167, 170, 179, 194, 199, 202, 215, 217, 230, 231, 233, 237, 239, 241, 243, 246, 251, 261, 269, 277, 280, 284, 323, 330, 331, 332, 337, 338, 346, 350, 357, 358, 359, 381, 449, 479, 481, 485, 494, 500, 502, 510, 511, 519, 520, 523, 524, 533, 534, 539, 540, 542, 543, 544, 545, 549, 561, 562, 563, 564, 565, 566, 568, 570, 575, 576, 578, 580, 581, 582, 585, 586, 589, 593, 595, 597, 598, 599, 601, 603, 604, 606, 607, 608, 609, 610, 612, 615, 616, 618, 619, 622, 623, 626, 628, 629, 631, 632, 635, 636, 638, 639, 641, 642, 643, 644, 645, 646, 648, 649, 650, 652, 653, 654, 656, 657, 659, 661,

663, 664, 667, 668, 669, 671, 672, 673, 676, 679, 687, 688, 692, 693, 695, 697, 698, 703, 708, 712, 713, 718, 719, 721, 722, 726, 727, 729, 730, 731, 732, 736, 737, 740, 741, 743, 744, 745, 746, 747, 748, 749, 750, 753, 755, 756, 758, 759, 760, 761, 763, 767, 768, 769, 771, 772, 773, 776, 777, 780, 781, 783, 785, 786, 787, 788, 789, 790, 791, 792, 793, 794, 795, 796, 797, 798, 799, 800, 801, 802, 803, 807, 811, 812, 856, 857, 859, 861, 863, 866, 868, 869, 871, 873, 874, 875, 876, 877, 878, 879, 880, 881, 882, 884, 885, 886, 888, 890, 892, 893, 894, 895, 899, 900, 909, 911, 913, 914, 915, 916, 917, 918, 919, 920, 921, 922, 925, 926, 928, 930, 934, 935, 936, 937, 951, 954, 962, 964, 965, 966, 972, 979, 980, 994, 995, 998, 1002, 1007, 1015, 1018, 1024, 1025, 1028, 1029, 1036, 1037, 1039, 1041, 1043, 1044, 1046, 1047, 1048, 1049, 1052, 1053, 1061, 1063, 1074, 1075, 1082, 1083, 1084, 1090, 1091, 1097, 1098, 1102, 1103, 1104, 1107, 1108, 1109, 1114, 1115, 1134, 1166, 1171, 1189, 1215, 1225, 1241, 1249  
 identified for deletion processing with subtables • 870, 909  
 identified for deletion processing without subtables • 869, 909  
 identified for insertion of source table • 877, 884  
 identified for replacement processing with subtables • 882, 890, 912  
 identified for replacement processing without subtables • 882, 890, 911  
 <identifier> • **6**, 45, 81, 87, 127, **161**, 162, 163, 164, 168, 189, 198, 201, 327, 328, 329, 330, 366, 367, 368, 369, 370, 374, 421, 474, 667, 675, 755, 962, 993, 1012, 1036, 1127, 1129, 1138, 1139, 1215, 1229  
 <identifier body> • **143**, 150  
 <identifier chain> • 90, **198**, 200, 1215  
 <identifier extend> • **143**, 148  
 <identifier part> • **143**, 149, 151, 1171, 1260  
 <identifier start> • **143**, 148, 151, 1171  
 identify • **6**  
 IDENTITY • 147, 574, 575, 576, 584  
 identity column • 60, 587, 591, 606, 874, 879, 919, 925  
 <identity column specification> • 580, **584**, 586, 587, 588, 604, 1211  
 <identity option> • **574**, 575  
 IMMEDIATE • 146, 545, 546, 593, 638, 661, 951, 955, 1034, 1181  
 immediate • 67, 68, 93, 127, 545, 593, 639, 836, 951, 1121  
 immediately dependent • 563  
 immediately subordinate descriptor areas • 983, 984, 1000  
 immediately subordinate descriptors • 983  
 impacted • 635, 797, 800, 802  
 impacted dereference operation • 610, 611, 622, 635  
 IMPLEMENTATION • 146, 526, 717, 725, 737, 1236  
 implementation-defined • **6**  
 <implementation-defined character set name> • **539**, 540



- implementation-defined classes* • 1143
- implementation-defined exception code* • 1228
- implementation-defined subclasses* • 1143
- implementation-dependent • **6**
- implicitly invocable • 50, 51, 746
- <implicitly typed value specification> • **196**, 215, 589
- IN • 147, 256, 257, 273, 274, 395, 412, 679, 698, 704, 715, 723, 728, 729, 730, 1005, 1012, 1022
- <in predicate> • 251, 372, 401, 402, **412**, 413, 478, 1009, 1177, 1273
- <in predicate part 2> • 211, **412**
- <in predicate value> • **412**
- in usage by • 638, 644
- <in value list> • **412**, 413, 1177, 1273
- <in-line window specification> • 61, **207**, 209, 369, 1245
- inappropriate access mode for branch transaction* • 948, 1149
- inappropriate isolation level for branch transaction* • 948, 1149
- INCLUDING • 146, 574, 575, 576
- <inclusive user-defined type specification> • **449**
- inclusively specified • 449
- INCREMENT • 146, 576, 766
- increment • 60, 82, 83, 504, 506, 507, 508, 575
- independent • 71, 103, 333, 790, 869, 871, 876, 881, 889, 908, 911, 1239, 1240
- independent node • 789
- <independent variable expression> • 65, 66, 205, 547, **548**, 550, 553, 554, 557, 558, 1174, 1175
- INDICATOR • 147, 191, 996, 998, 1000, 1019, 1024, 1030, 1250
- indicator overflow* • 455, 1145
- <indicator parameter> • **191**, 193, 194, 371, 1177, 1231
- <indicator variable> • **191**, 192, 194, 371, 1177, 1231
- inherited column • 578
- inherited columns • 59
- inherited method • 41
- initial cursor position • 94
- INITIALLY • 146, 545, 546, 593, 638, 661, 1181
- INNER • 74, 75, 76, 147, 335, 336, 338, 341
- innermost executing statement • 129
- INOUT • 147, 715, 723, 728, 730, 1012
- INPUT • 89, 146, 677, 680, 698, 717, 720, 1017
- <input using clause> • **1023**, 1026, 1032, 1039, 1040, 1161, 1250
- INSENSITIVE • 94, 100, 102, 147, 851, 852, 854, 855, 903, 1183
- insensitive • 102, 103, 869, 871, 876, 881, 889, 903, 908, 910
- INSERT • 69, 120, 121, 134, 135, 136, 147, 561, 562, 564, 565, 582, 605, 617, 666, 667, 668, 670, 671, 672, 777, 779, 780, 781, 786, 787, 794, 798, 873, 875, 877, 878, 880, 881, 898, 920, 923, 940, 1168, 1181, 1237
- <insert column list> • 794, 798, **873**, 874, 875, 876, 877, 878, 879, 880, 884, 885, 1008, 1189
- <insert columns and source> • **873**, 874, 875, 1008
- insert INSTEAD OF trigger • 133
- insert object column • 880
- <insert statement> • 51, 60, 107, 110, 114, 115, 116, 133, 391, 582, 794, 795, 797, 798, 799, 832, **873**, 874, 875, 877, 1003, 1008, 1119, 1132, 1134, 1135, 1178, 1183, 1209, 1251, 1264
- insert trigger • 133
- insertable-into • 55, 56, 57, 331, 373, 381, 564, 581, 626, 631, 667, 670, 671, 873, 874, 879, 880, 923
- insertable-into derived table • 331
- <insertion target> • **873**, 877, 1209
- INSTANCE • 146, 541, 543, 676, 679, 689, 697, 708, 716, 1188
- instance • **6**
- instance SQL-invoked methods • 88
- INSTANTIABLE • 146, 674, 676, 677, 684, 687, 688, 1188
- <instantiable clause> • **674**, 676, 677, 688, 689, 1188, 1190
- INSTEAD • 54, 55, 133, 134, 146, 666, 667, 668, 669, 670, 671, 672, 917, 923, 930, 1213
- INSTEAD OF trigger • 133
- insufficient item descriptor areas* • 1019, 1151
- INT • 30, 147, 174, 176, 178, 182, 815, 820, 822, 825, 828, 1074, 1075, 1076, 1100, 1101, 1230
- INTEGER • 14, 15, 30, 60, 147, 174, 176, 178, 182, 409, 469, 474, 524, 726, 729, 730, 750, 822, 840, 841, 843, 844, 845, 847, 984, 985, 1011, 1076, 1085, 1092, 1094, 1098, 1101, 1105, 1109, 1116, 1230, 1287
- integrity constraint violation* • 229, 836, 934, 935, 936, 937, 955, 1068, 1136, 1139, 1148, 1150
- interface • **8**
- intermediate result table • 388, 389
- internal sequence generator • 82
- INTERSECT • 23, 29, 50, 79, 80, 147, 251, 252, 308, 376, 379, 380, 381, 382, 384, 385, 390, 391, 408, 481, 483, 1169, 1219, 1248
- INTERSECTION • 147, 252, 310, 483, 484, 547, 550, 553, 558, 1201, 1202
- INTERVAL • 14, 15, 34, 35, 98, 129, 147, 154, 158, 175, 183, 289, 290, 293, 294, 295, 296, 409, 439, 440, 470, 475, 840, 842, 843, 845, 846, 847, 973, 1001, 1002, 1005, 1006, 1010, 1011
- <interval absolute value function> • 39, **298**
- <interval factor> • **293**, 294
- interval field overflow* • 228, 297, 457, 462, 1139, 1145

*<interval fractional seconds precision>* • 38, 159, 182, 260, **509**, 510, 511, 984, 1006, 1021, 1235  
*<interval leading field precision>* • 38, 181, 293, 294, 295, 296, 409, **509**, 510, 511, 512, 984, 1006, 1021, 1233, 1235  
*<interval literal>* • 152, **154**, 158, 159, 160, 590, 1166  
*<interval primary>* • 288, 289, 290, **293**, 294, 1233  
*<interval qualifier>* • 34, 38, 154, 158, 159, 175, 179, 182, 183, 228, 293, 294, 295, 296, 297, 498, **509**, 510, 511, 512, 590, 838, 984, 991, 1005, 1021, 1166, 1235  
*<interval string>* • 144, 148, **154**  
*<interval term>* • 288, 289, 290, **293**, 294  
*<interval term 1>* • **293**, 294, 295, 1233  
*<interval term 2>* • **293**, 295  
*<interval type>* • 15, 34, 173, **175**, 179, 183, 184, 358, 470, 498, 613, 984, 1006, 1011, 1021, 1166, 1232  
*<interval value expression>* • 250, 251, 252, 257, 259, 260, 261, 288, 289, 290, **293**, 294, 295, 296, 297, 298, 973, 1010, 1166, 1247  
*<interval value expression 1>* • **293**, 294, 295, 1233  
*<interval value function>* • 293, 294, **298**, 1166  
*interval value out of range* • 551, 1145  
*intervals* • 15, 37, 38, 39, 158, 180, 228, 409, 1247  
*INTO* • 147, 209, 582, 857, 862, 873, 878, 1027  
*<into argument>* • **1027**, 1028, 1031, 1241  
*<into arguments>* • **1027**, 1028, 1031  
*<into descriptor>* • **1027**, 1028, 1030  
*<introducer>* • **152**, 156, 157  
*invalid* • 97, 129, 290, 532, 534, 826, 899, 956, 957, 958, 1011, 1048, 1049, 1050, 1052, 1053, 1137  
*invalid argument for natural logarithm* • 262, 1145  
*invalid argument for power function* • 263, 1145  
*invalid argument for width bucket function* • 264, 1145  
*invalid authorization specification* • 813, 961, 962, 971, 1148  
*invalid catalog name* • 974, 1148  
*invalid character set name* • 977, 1148  
*invalid character value for cast* • 220, 221, 222, 223, 224, 229, 1139, 1145  
*invalid collation name* • 980, 1148  
*invalid condition number* • 945, 948, 1136, 1148  
*invalid connection name* • 962, 1148  
*invalid cursor name* • 1011, 1018, 1036, 1039, 1041, 1043, 1044, 1046, 1148  
*invalid cursor state* • 858, 901, 905, 907, 909, 911, 1015, 1137, 1148  
*invalid DATA target* • 1000, 1147  
*invalid datetime format* • 224, 225, 226, 227, 1145  
*invalid DATETIME\_INTERVAL\_CODE* • 1002, 1147  
*invalid descriptor count* • 1024, 1028, 1147  
*invalid descriptor index* • 993, 997, 1000, 1147  
*invalid escape character* • 282, 416, 422, 1145  
*invalid escape octet* • 417, 1145  
*invalid escape sequence* • 416, 417, 1145  
*invalid grantor* • 787, 804, 1148  
*invalid indicator parameter value* • 459, 1145  
*invalid interval format* • 228, 1145  
*invalid LEVEL value* • 1000, 1001, 1147  
*invalid number of conditions* • 948, 1151  
*invalid parameter value* • 813, 1145  
*invalid preceding or following size in window function* • 361, 362, 363, 364, 1145  
*invalid regular expression* • 282, 422, 1145  
*invalid repeat argument in a sample clause* • 333, 1146  
*invalid role specification* • 783, 972, 1148  
*invalid sample size* • 333, 1146  
*invalid schema name* • 975, 1148  
*invalid schema name list specification* • 978, 1148  
*invalid specification* • 532, 534, 826, 899, 900, 954, 958, 1149, 1150  
*invalid SQL descriptor name* • 993, 994, 995, 997, 1000, 1018, 1023, 1028, 1148  
*invalid SQL statement identifier* • 1012, 1149  
*invalid SQL statement name* • 1015, 1018, 1032, 1037, 1039, 1149, 1255  
*invalid SQL-invoked procedure reference* • 1037, 1148  
*invalid target type specification* • 234, 1149  
*invalid time zone displacement value* • 290, 973, 1146  
*invalid transaction state* • 834, 835, 836, 869, 871, 876, 881, 889, 908, 910, 945, 947, 948, 966, 971, 972, 1121, 1122, 1149  
*invalid transaction termination* • 955, 957, 1149  
*invalid transform group name specification* • 979, 1149  
*invalid use of escape character* • 282, 422, 1146  
*invalid XQuery option flag* • 493, 1146  
*invalid XQuery regular expression* • 493, 1146  
*invalid XQuery replacement string* • 496, 1146  
*inverse distribution* • 64, 66  
*<inverse distribution function>* • **548**, 550, 551, 555, 558, 1175, 1221, 1236  
*<inverse distribution function argument>* • **548**, 555, 558, 1175  
*<inverse distribution function type>* • **548**  
*invocable routine* • 517, 518  
*INVOKER* • 146, 180, 203, 217, 218, 244, 280, 332, 540, 544, 630, 717, 732, 733, 734, 735, 737, 867, 868, 875, 880, 887, 888, 1044, 1046, 1049, 1053  
*IS* • 52, 147, 212, 270, 283, 299, 300, 301, 302, 308, 309, 311, 428, 434, 441, 448, 449, 450, 573, 586, 595, 675, 1071, 1072, 1073, 1074, 1075, 1078, 1079, 1080, 1081,

1082, 1083, 1084, 1087, 1088, 1089, 1090, 1091, 1094, 1095, 1096, 1097, 1098, 1100, 1101, 1102, 1103, 1105, 1106, 1107, 1108, 1109, 1111, 1112, 1113, 1114, 1115  
 is destroyed • 88, 100, 123, 135, 526, 537, 538, 572, 617, 621, 624, 637, 645, 647, 651, 655, 659, 664, 673, 696, 711, 714, 744, 748, 769, 785, 907, 941, 953, 995, 1015, 1037  
 is established • 123  
 ISOLATION • 146, 945, 949  
 <isolation level> • 124, 125, 126, 127, 131, 524, 537, 945, 947, 948, **949**, 950, 956, 958, 962, 963, 969, 1121, 1149, 1167, 1240, 1275  
 <item number> • **996**, 997, 998, 999, 1000  
 iteration ignorable • 388

## — J —

JOIN • 75, 147, 335, 338, 1134, 1135  
 <join column list> • 335, **336**, 337  
 <join condition> • 75, 329, 330, **335**, 336, 338, 339, 384, 1268  
 join partitioning column • 336, 337, 338  
 <join specification> • 74, 75, **335**, 336, 338  
 <join type> • 75, **335**, 336, 379, 380, 1268  
 <joined table> • v, 23, 29, 74, 77, 199, 325, 326, 329, 330, 331, 332, 333, **335**, 336, 337, 338, 339, 343, 368, 378, 478, 479, 879, 1172, 1233, 1268

## — K —

K • 143, 146, 176  
 KEY • 52, 68, 73, 146, 578, 586, 595, 596, 597, 598, 618  
 <key word> • 14, 18, 143, **146**, 150, 164, 299, 300, 987, 989, 997, 999, 1225, 1243  
 key-preserving • 372, 373  
 key-preserving with respect to • 372  
 KEY\_MEMBER • 146, 996, 998, 999, 1020, 1215  
 KEY\_TYPE • 146, 996, 999, 1018, 1019, 1020  
 KIND • 1094  
 known functional dependencies • 72, 74, 75, 76, 77, 78, 79, 80  
 known functional dependency • 72, 73, 74, 75, 76, 77, 78, 79  
 known not null • 72, 73, 74, 338, 371, 863, 1226, 1233  
 known not nullable • 52, 58, 74, 75, 382, 383, 575, 577, 618, 630, 644, 662, 664  
 known-not-null condition • 300, 1233

## — L —

LANGUAGE • 147, 513, 677, 679, 680, 686, 687, 688, 698, 699, 720  
 <language clause> • 40, 87, 89, **513**, 514, 676, 680, 685, 698, 716, 720, 731, 734, 735, 738, 739, 740, 741, 742,

805, 806, 808, 809, 810, 811, 812, 1021, 1059, 1163, 1164, 1165  
 <language name> • 42, **513**, 685, 686, 701, 706, 734, 735, 741, 742, 812  
 LARGE • 14, 15, 18, 28, 98, 147, 173, 174, 176, 177, 178, 182, 229, 266, 419, 425, 453, 454, 468, 469, 474, 585, 826, 827, 829, 838, 839, 840, 841, 843, 844, 845, 846, 984, 985, 989, 990, 1001, 1005, 1031, 1084, 1173, 1206, 1207  
 large object character string • 18, 96, 216, 222, 223, 275, 532, 533, 826  
 <large object length> • **174**, 176, 177, 178, 838, 1072, 1080, 1088, 1095, 1101, 1105, 1112  
 <large object length token> • **143**, 149, 174, 176  
 large object string types • 15  
 large object strings • 15, 29, 269, 271  
 LAST • 61, 62, 146, 361, 362, 363, 559, 560, 857, 905, 906  
 LATERAL • 147, 325, 329, 333, 1218  
 <lateral derived table> • 76, **325**, 329, 330, 331, 333, 336, 340, 1218  
 LEADING • 147, 274, 285, 287, 1088  
 leaf column • 215  
 leaf generally underlying table • 57  
 leaf table • 59  
 leaf underlying table • 57  
 least significant • 510  
 LEFT • 74, 75, 76, 147, 335, 336, 338, 339, 341, 379, 380  
 <left brace> • 23, 140, **141**, 420, 422  
 <left bracket> • 23, 140, **141**, 421, 422, 424, 1078, 1105  
 <left bracket or trigraph> • **141**, 175, 191, 196, 248, 306, 312, 892  
 <left bracket trigraph> • **141**, 145  
 <left paren> • 23, 139, **140**, 173, 174, 175, 188, 191, 200, 205, 207, 211, 215, 233, 235, 237, 246, 249, 256, 257, 258, 273, 274, 275, 291, 293, 298, 299, 306, 311, 315, 325, 326, 335, 345, 356, 366, 376, 381, 398, 412, 420, 422, 424, 449, 509, 516, 541, 547, 548, 573, 574, 584, 595, 597, 601, 625, 631, 661, 666, 674, 675, 715, 716, 745, 747, 755, 758, 759, 761, 779, 811, 866, 873, 878, 892, 1004, 1005, 1055, 1057, 1071, 1072, 1079, 1080, 1087, 1088, 1094, 1095, 1100, 1101, 1105, 1111, 1112, 1261, 1263  
 LENGTH • 146, 822, 829, 984, 985, 996, 1000, 1001, 1002, 1020, 1021, 1074, 1241, 1250  
 <length> • 85, **174**, 176, 177, 178, 180, 279, 838, 1072, 1073, 1079, 1081, 1088, 1089, 1094, 1096, 1101, 1102, 1105, 1106, 1107, 1111, 1113, 1115  
 <length expression> • 22, 29, 32, 256, **257**, 260, 266, 1173, 1232  
 length in positions • 34, 511, 512, 1021  
 <less than operator> • 140, **141**, 403, 560

<less than or equals operator> • 144, **145**, 403, 404  
 LEVEL • 88, 146, 394, 716, 720, 734, 945, 949, 983, 994, 996, 1000, 1001, 1019, 1020, 1021, 1024, 1028  
 <level of isolation> • 947, 948, **949**, 950, 969, 1167, 1240  
 <levels clause> • **625**, 626, 633, 1181  
 LIKE • 22, 147, 414, 415, 416, 417, 418, 419, 425, 574, 582, 583, 1169, 1211  
 <like clause> • 573, **574**, 575, 580, 582, 583, 897, 1211  
 <like option> • **574**  
 <like options> • **574**, 575, 583, 1211  
 <like predicate> • 22, 23, 29, 401, 402, **414**, 415, 417, 419, 478, 1173, 1207, 1234, 1261  
 LIKE\_REGEX • 23, 147, 426, 1255  
 linearly recursive • 378  
 LIST • 265  
 <list of attributes> • **675**, 678, 685  
 <literal> • 75, 77, **152**, 168, 171, 190, 193, 220, 222, 223, 224, 225, 226, 227, 228, 351, 367, 556, 589, 590, 591, 893, 901, 902, 994, 999, 1004, 1123, 1139, 1162, 1282  
 LN • 147, 258, 263  
 LOB-ordered • 17, 343, 479, 484, 1172, 1207  
 LOCAL • 55, 57, 100, 127, 147, 164, 168, 170, 288, 289, 296, 573, 580, 625, 627, 631, 853, 897, 924, 931, 947, 948, 973, 1050, 1213  
 local • 36, 54, 55, 56, 57, 87, 99, 129, 202, 524, 525, 531, 537, 570, 573, 580, 581, 582, 598, 601, 603, 622, 626, 667, 725, 780, 897, 898, 1137, 1176, 1228, 1235, 1245  
 <local or schema qualified name> • **161**, 164, 167  
 <local or schema qualifier> • 161, **162**, 164, 171, 575, 897, 1184  
 <local qualified name> • **162**  
 <local qualifier> • **162**, 165, 171, 1184, 1220  
 local time • 35, 36, 160  
 LOCALTIME • 147, 291  
 LOCALTIMESTAMP • 147, 291, 292, 301, 602, 662, 1179  
 LOCATOR • 96, 146, 675, 679, 681, 697, 700, 715, 899, 900, 984, 985, 986, 989, 990, 998, 1000, 1001, 1024, 1030, 1031, 1060, 1062, 1072, 1074, 1075, 1080, 1083, 1084, 1088, 1090, 1091, 1095, 1097, 1098, 1101, 1102, 1103, 1105, 1106, 1108, 1112, 1114, 1115  
 locator • 92, 93, 96, 97, 129, 453, 454, 524, 532, 534, 706, 734, 807, 811, 812, 826, 829, 899, 900, 956, 957, 958, 986, 1025, 1028, 1029, 1030, 1058, 1060, 1061, 1062, 1063, 1074, 1075, 1083, 1090, 1091, 1097, 1098, 1102, 1103, 1108, 1114, 1115, 1283  
 locator exception • 532, 534, 826, 899, 900, 1149  
 <locator indication> • 92, 93, 679, 680, 681, 682, 683, 688, 689, 698, 699, 701, 703, 704, **715**, 716, 717, 721, 722, 723, 726, 727, 728, 729, 730, 734, 735, 737, 738, 806, 811, 813, 1196, 1197, 1198, 1205, 1219, 1220, 1273, 1283  
 <locator reference> • **899**, 900, 1011

<low value> • **420**, 422, 423  
 LOWER • 22, 147, 273, 283, 421, 424

## — M —

M • 143, 146, 176, 489, 513, 514, 529, 731, 739, 741, 809, 824, 825, 828, 1059, 1064, 1102, 1164  
 <major category> • **1067**, 1068  
 <mantissa> • 31, **153**, 159, 221, 223  
 MAP • 17, 41, 45, 146, 404, 406, 442, 478, 479, 485, 687, 749, 750, 751, 752, 1200  
 <map category> • **749**, 750  
 <map function specification> • 17, 479, 485, **749**, 750  
 marked modified privilege descriptor • 790  
 MATCH • 68, 69, 147, 436, 438, 597, 599, 600, 1181  
 match • 90, 265, 266, 282, 284, 285, 426, 493, 494, 495, 985, 997, 1000, 1023, 1024, 1027, 1028, 1147  
 <match predicate> • 69, 251, 401, 402, **436**, 437, 438, 478, 479, 1009, 1181  
 <match predicate part 2> • 212, **436**  
 <match type> • 68, 69, **597**, 600, 933  
 MATCHED • 146, 878  
 matching row • 69, 70, 882, 933, 934, 935, 936, 937  
 MAX • 65, 147, 252, 408, 485, 547, 549, 552, 556, 558, 993, 1208, 1241, 1248, 1249  
 maximal supertable • 58  
 <maximum cardinality> • **175**, 179, 184, 1194  
 maximum number of stacked diagnostics areas exceeded • 1142, 1147  
 <maximum returned result sets> • 716, **717**, 733, 741  
 maximum value • 60  
 MAXVALUE • 146, 505, 507, 576, 766, 1235  
 MEMBER • 147, 444  
 <member> • **674**  
 <member list> • 179, **674**, 677, 678, 687, 688, 1187  
 <member name> • **541**  
 <member name alternatives> • **541**  
 <member predicate> • 251, 372, 401, 402, **444**, 445, 478, 479, 483, 1201  
 <member predicate part 2> • 212, **444**  
 MERGE • 147, 878, 885, 1169  
 <merge correlation name> • **878**, 879, 883, 1134  
 <merge insert specification> • 794, 798, **878**, 1008  
 <merge insert value element> • **878**, 879, 880  
 <merge insert value list> • 230, 329, **878**, 879, 880, 1008  
 <merge operation specification> • **878**, 881, 882, 883  
 <merge statement> • 51, 60, 107, 111, 114, 115, 117, 133, 329, 794, 795, 796, 797, 798, 799, 832, **878**, 879, 880, 881, 883, 884, 885, 892, 909, 911, 1003, 1008, 1119, 1132, 1134, 1135, 1169, 1170, 1183, 1189, 1209, 1250, 1251

- <merge update specification> • **878**
- <merge when clause> • **878**, 881
- <merge when matched clause> • **878**, 879, 880, 881, 882, 885, 1134, 1135, 1170
- <merge when not matched clause> • **878**, 879, 880, 883, 885, 1134, 1135, 1170
- message text item • 528, 529
- MESSAGE\_LENGTH • 146, 1126, 1138
- MESSAGE\_OCTET\_LENGTH • 146, 1126, 1138
- MESSAGE\_TEXT • 146, 1126, 1138, 1243
- METHOD • 147, 541, 542, 543, 676, 677, 708, 716, 738, 1188, 1190
- <method characteristic> • **88**, **676**
- <method characteristics> • **675**, **676**, 680, 688, 697, 1188
- <method invocation> • **89**, 188, 189, 200, **235**, 236, 239, 244, 372, 695, 713, 740, 743, 792, 793, 794, 796, 797, 1187
- <method name> • 40, 41, 42, 89, **162**, 169, 235, 237, 244, 541, 542, 543, 632, 676, 678, 680, 682, 683, 685, 686, 697, 701, 704, 706, 708, 710, 716, 719, 720, 771, 773, 892, 893, 894, 896, 1189
- method of type • 41
- <method reference> • 241, **244**, 245, 372, 610, 622, 632, 635, 695, 713, 740, 743, 771, 773, 792, 793, 794, 795, 796, 797, 798, 1191
- <method selection> • **235**, 236, 516, 518, 519
- <method specification> • 40, 42, 88, **675**, 678, 688, 689, 1219, 1220
- method specification descriptor • 42, 93, 517, 624, 636, 654, 679, 685, 698, 703, 706, 708, 709, 718, 719, 721, 724
- <method specification designator> • 88, 715, **716**, 718, 719, 732, 733, 735, 737, 738, 1187, 1190
- <method specification list> • 40, 42, 674, **675**, 678, 685, 686, 689, 1187
- MIN • 65, 147, 252, 408, 485, 547, 549, 552, 556, 558, 1208, 1248, 1249
- Minimal common supertype • 465
- minimal common supertype • 44, 465
- minimum value • 60, 65, 82, 83, 157, 504, 506, 507, 508, 552, 576
- <minus sign> • 23, 139, **140**, 145, 149, 153, 154, 254, 255, 288, 290, 293, 294, 421, 422, 423, 511, 512, 1006, 1166
- MINUTE • 35, 129, 147, 158, 181, 261, 289, 464, 509, 510, 973, 1002, 1010
- <minutes value> • **154**, 158
- MINVALUE • 146, 505, 507, 576, 766, 1235
- MOD • 147, 257, 266, 1218
- modified • 81, 88, 91, 102, 125, 508, 603, 606, 614, 641, 646, 655, 670, 671, 672, 692, 694, 741, 742, 758, 789, 790, 800, 801, 802, 938, 1073, 1075, 1081, 1084, 1089, 1091, 1095, 1098, 1101, 1103, 1106, 1109, 1112, 1115, 1136
- MODIFIES • 147, 717, 720, 732, 736, 742
- modifying SQL-data • 130
- modifying SQL-data not permitted* • 525, 527, 531, 1148, 1150
- MODULE • 55, 147, 162, 164, 165, 167, 202, 203, 810, 897, 898, 1137, 1184
- <module authorization clause> • 85, 167, **805**, 806, 808, 1060, 1163, 1229, 1241
- <module authorization identifier> • 83, 84, 119, 568, 797, **805**, 806, 812, 813, 961, 1013, 1056, 1240, 1251
- <module character set specification> • 84, 797, **810**, 1059, 1175, 1238
- <module collation specification> • **805**, 806, 807, 808, 1180
- <module collations> • **805**, 1056, 1060
- <module contents> • 99, 164, 165, **805**
- <module name clause> • **805**, **810**, 1059, 1175
- <module path specification> • 91, 736, **805**, 806, 808, 1060, 1192, 1238, 1242
- <module transform group specification> • **805**, 806, 808, 1060, 1199
- <modulus expression> • 32, 256, **257**, 260, 262, 266, 1005, 1218
- monadic • **8**
- MONTH • 35, 37, 147, 158, 179, 181, 361, 464, 509, 510, 1002, 1006
- <months value> • **154**, 158, 511
- MORE • 146, 537, 1125, 1129
- more recent execution context • 113
- most recent atomic execution context • 113
- most recent execution context • 113
- most significant • 510
- most specific type • **8**
- most specific type mismatch* • 535, 1146
- <multiple column assignment> • **892**, 893, 896, 1010, 1222
- <multiple group specification> • **717**, 722, 806, 1058
- multiple server transactions* • 947, 961, 964, 1148
- <multiplier> • **143**, 149, 151, 174, 176, 1207
- MULTISET • 13, 47, 48, 49, 50, 99, 147, 175, 183, 196, 197, 219, 251, 252, 308, 309, 310, 311, 312, 470, 483, 498, 549, 589, 718, 840, 842, 844, 845, 846, 847, 875, 880, 894, 985, 986, 998, 1000, 1007, 1024, 1030, 1031, 1200, 1201
- multiset • 826
- <multiset element> • **312**, 313, 1007
- multiset element grouping operation • 483
- <multiset element list> • **312**, 313, 1007
- <multiset element reference> • 50, 188, 189, **249**, 372, 1200
- multiset locator parameter • 812, 826, 829, 899, 900

multiset locator variable • 1075, 1083, 1091, 1098, 1103, 1108, 1115

multiset operand • 479, 480, 483, 484, 1188, 1202

<multiset primary> • **308**, 309, 483

<multiset set function> • 50, 252, **311**, 483

<multiset term> • **308**, 309, 310, 483

<multiset type> • 13, 48, 97, **175**, 180, 184, 219, 251, 308, 311, 312, 372, 442, 446, 448, 454, 457, 458, 462, 465, 479, 480, 483, 532, 533, 534, 538, 550, 589, 680, 681, 689, 698, 699, 721, 722, 737, 738, 811, 812, 813, 875, 880, 894, 1029, 1062, 1072, 1075, 1080, 1083, 1088, 1091, 1095, 1098, 1101, 1103, 1106, 1108, 1112, 1115, 1195, 1198, 1200, 1201, 1202, 1220

<multiset value constructor> • 188, 189, **312**, 313, 1201, 1215

<multiset value constructor by enumeration> • **312**, 313

<multiset value constructor by query> • **312**, 313

<multiset value expression> • 249, 250, 251, 252, **308**, 309, 310, 311, 444, 446, 483, 1201

<multiset value function> • 308, 309, **311**, 1201

*multiset value overflow* • 551, 1146

multiset-ordered • 17, 343, 479, 481, 484, 485, 1172, 1202

MUMPS • 146, 513, 739, 809, 1104, 1153, 1159, 1164

<MUMPS array locator variable> • 1100, **1101**, 1103, 1104, 1197

<MUMPS BLOB locator variable> • 1100, **1101**, 1102, 1104, 1206

<MUMPS BLOB variable> • 1100, **1101**, 1102, 1104, 1206

<MUMPS character variable> • **1100**, 1101

<MUMPS CLOB locator variable> • 1100, **1101**, 1103, 1104, 1206

<MUMPS CLOB variable> • **1100**, 1101, 1102, 1104, 1206

<MUMPS derived type specification> • **1100**, 1101

<MUMPS host identifier> • 1057, **1100**, 1101, 1102, 1103

<MUMPS length specification> • **1100**, 1101

<MUMPS multiset locator variable> • 1100, **1101**, 1103, 1104, 1198

<MUMPS numeric variable> • **1100**

<MUMPS REF variable> • 1100, **1101**, 1103, 1104, 1191

<MUMPS type specification> • **1100**

<MUMPS user-defined type locator variable> • 1100, **1101**, 1103, 1104, 1196

<MUMPS user-defined type variable> • 1100, **1101**, 1102, 1104, 1199

<MUMPS VARBINARY variable> • **1100**, 1102, 1104, 1203

<MUMPS variable definition> • 1056, **1100**, 1101, 1103

<mutated set clause> • **892**, 893, 894, 896, 1010, 1189

<mutated target> • **892**, 893, 894

mutator function • **8**

mutually recursive • 378

## — N —

N • 152

n-adic operator • **8**

NAME • 146, 717, 740, 996, 999, 1019, 1020, 1250

<named columns join> • 74, 75, **335**, 336, 337, 338, 339, 342, 343, 1172

NAMES • 146, 810, 977, 1056

NATIONAL • 18, 147, 173, 174, 176, 229, 266, 419, 821, 822, 1084, 1098, 1173, 1206, 1207, 1225, 1230

<national character large object type> • **174**, 184, 1204

<national character string literal> • 143, 148, **152**, 155, 156, 160, 1173

<national character string type> • **173**, 184, 1173

NATURAL • 74, 75, 147, 335, 336, 338, 339, 342, 478, 479, 616

<natural join> • **335**, 336, 343, 379, 1172

<natural logarithm> • 32, 256, **258**, 260, 262, 266, 1221, 1232

NCHAR • 147, 173, 174, 176, 815, 821, 822, 1079, 1081, 1084

NCLOB • 147, 174, 176, 1079, 1081, 1084, 1088, 1089

negative • 82, 159, 181, 182, 222, 262, 263, 290, 294, 301, 361, 362, 363, 364, 490, 505, 531, 532, 534, 906, 998, 1024, 1129, 1134

NESTING • 146, 1017, 1018, 1019

<nesting option> • **1017**

NEVER • 587

NEW • 147, 239, 666, 667, 668, 716, 720, 734

<new invocation> • 235, 236, **239**, 522

new row • 135

new savepoint level • 88, 93, 123, 538, 941, 956, 957, 963

<new specification> • 89, 188, 189, **239**, 240, 1187

new transition table • 133, 866, 1048

<new transition table name> • 135, **666**, 667, 669

new transition variable • 134

new transition variable column reference • 200

<new transition variable name> • 136, 200, 666, **667**, 669

<new window name> • **356**, 357, 359

<newline> • 145, **146**, 148, 149, 156, 1229

NEXT • 146, 230, 857, 905, 906, 1041

<next value expression> • 83, 188, 189, **230**, 231, 251, 1211

NFC • 146, 274, 279, 434

NFD • 146, 274

NFKC • 146, 274

NFKD • 146, 274

niladic • **8**

NO • 19, 25, 69, 70, 83, 94, 100, 147, 357, 365, 408, 504, 505, 506, 507, 540, 574, 597, 599, 652, 680, 688, 698,

- 717, 720, 732, 742, 766, 767, 836, 851, 904, 934, 935, 936, 937, 955, 957, 980, 1038, 1188, 1235
- no active SQL-transaction for branch transaction* • 947, 1149
- no additional result sets returned* • 1037, 1149
- no data* • 95, 96, 490, 491, 836, 837, 858, 863, 870, 872, 877, 884, 890, 906, 997, 1037, 1041, 1042, 1122, 1123, 1143, 1149
- no subclass* • 1143, 1144, 1147, 1148, 1149, 1150, 1151
- non-aggregated argument • 205
- non-atomic • 112, 113, 132, 837
- non-atomic execution context • 113
- non-atomic SQL statement • 112
- <non-cycle mark value> • **393**, 394
- non-deferrable • 67
- <non-escaped character> • **420**, 422
- <non-extended descriptor name> • 163, **164**, 168, 170, 171, 993, 1162, 1163
- non-extended names • 87
- non-linearly recursive • 379
- non-recursive • 377, 383, 387
- non-recursive operand • 378
- <non-reserved word> • **146**, 1255
- <non-second primary datetime field> • **509**, 511
- non-unique matching row • 933
- noncharacter • **6**
- noncharacter in UCS string* • 180, 1146
- <nondelimiter token> • 10, **143**, 149
- <nondoublequote character> • **144**, 148
- NONE • 41, 147, 685, 749, 754, 972
- nonholdable • 129
- <nonparenthesized value expression primary> • **188**, 189, 251, 299, 300, 301, 302, 318, 1204
- <nonquote character> • 145, **152**, 156
- <normal form> • 28, 30, 178, 179, 182, 183, **274**, 279, 287, 385, 386, 434, 435, 1172, 1230, 1277
- normalization • **6**
- NORMALIZE • 147, 158, 270, 274, 283
- <normalize function> • 28, 273, **274**, 275, 279, 281, 282, 287, 1208, 1226, 1233
- <normalize function result length> • **274**, 279, 1233
- NORMALIZED • 146, 270, 283, 434
- <normalized predicate> • 28, 401, **434**, 435, 1208, 1209, 1226
- <normalized predicate part 2> • 212, **434**
- NOT • 22, 23, 33, 147, 212, 299, 300, 301, 302, 404, 405, 411, 412, 414, 415, 420, 421, 426, 428, 434, 440, 441, 443, 444, 446, 448, 449, 450, 545, 546, 575, 576, 577, 578, 584, 586, 593, 595, 596, 638, 661, 674, 675, 676, 680, 688, 689, 698, 717, 720, 734, 878, 883, 1067, 1068, 1070, 1181, 1188, 1189, 1211, 1220
- <not equals operator> • 23, 29, 144, **145**, 403, 404, 410, 431, 478, 483, 485
- not permitted • 130
- not recursively referred to • 383
- not updatable • 52, 54, 101, 855
- NULL • 33, 52, 70, 89, 133, 147, 196, 212, 213, 218, 299, 300, 308, 309, 311, 327, 351, 372, 428, 459, 575, 576, 577, 578, 584, 586, 595, 596, 597, 599, 677, 680, 691, 698, 717, 720, 934, 935, 936, 937, 944, 1220, 1262
- <null ordering> • 62, 359, **559**, 560, 1220, 1237
- <null predicate> • 23, 401, 402, **428**, 429, 1272
- <null predicate part 2> • 211, **428**
- null row not permitted in table* • 320, 1146
- <null specification> • **196**, 197, 316, 589, 944, 1265
- null value • **6**
- null value eliminated in set function* • 552, 553, 555, 1151
- null value in array target* • 535, 859, 864, 895, 1146
- null value not allowed* • 193, 529, 1146, 1148
- null value substituted for mutator subject parameter* • 522, 691, 1146
- null value, no indicator parameter* • 454, 455, 998, 1146
- <null-call clause> • 89, 676, 680, 698, 716, **717**, 720, 740, 741, 742
- null-call function • 89, 526, 529, 531, 734, 736, 742
- nullability characteristic • 52, 58, 575, 577, 587, 618, 621, 630, 644, 645, 662, 664
- NULLABLE • 146, 996, 999, 1019, 1020
- NULLIF • 147, 211, 212
- NULLS • 61, 62, 146, 361, 362, 363, 559, 560
- NUMBER • 146, 1125, 1129
- <number of conditions> • 945, 947, 948, **949**, 950, 970, 1250
- numbers • 15, 30, 31, 32, 80, 148, 324, 340, 556, 1243
- NUMERIC • 14, 15, 30, 147, 174, 178, 182, 469, 474, 822, 840, 841, 843, 844, 845, 846, 984, 985, 1001, 1004, 1005, 1006, 1011, 1092, 1230
- <numeric primary> • **254**, 255, 1258
- <numeric type> • 173, **174**, 1257
- numeric types • 15, 469
- <numeric value expression> • 65, 248, 250, 252, **254**, 255, 257, 258, 260, 262, 263, 264, 275, 325, 548, 557, 1005, 1174, 1232, 1258
- <numeric value expression base> • **258**, 263
- <numeric value expression dividend> • 257, **258**, 262, 1005
- <numeric value expression divisor> • 257, **258**, 260, 262, 1005
- <numeric value expression exponent> • **258**, 263
- <numeric value function> • 254, **256**, 260, 1283
- numeric value out of range* • 220, 255, 262, 263, 264, 265, 456, 461, 551, 552, 555, 1139, 1146

## — O —

OBJECT • 14, 15, 18, 28, 98, 146, 173, 174, 176, 177, 178, 182, 229, 266, 419, 425, 453, 454, 468, 469, 474, 585, 826, 827, 829, 838, 839, 840, 841, 843, 844, 845, 846, 984, 985, 989, 990, 1001, 1005, 1031, 1084, 1173, 1206, 1207

object • 6

<object column> • 795, 798, 880, 882, 886, 887, 888, 889, 890, **892**, 893, 894, 895, 911, 1046, 1053

<object name> • 776, 777, **779**, 780, 781, 786, 802, 803, 1165, 1166, 1187, 1189, 1193

<object privileges> • 777, **779**, 780

object row • 882, 890, 911

observer function • 8

occupied • 46, 97, 98, 537, 837, 1084, 1122, 1136, 1139, 1243

OCCURRENCE • 256, 273, 274

<occurrences> • **993**, 994, 997, 1000, 1018, 1024, 1028, 1162, 1241

OCCURRENCES\_REGEX • 147, 256, 1255

<octet length expression> • **257**, 262, 1005, 1259, 1283

<octet like predicate> • **414**, 415, 417, 418, 1009, 1203, 1207

<octet like predicate part 2> • 211, **414**

<octet pattern> • **414**, 415, 478, 1009

OCTET\_LENGTH • 147, 257, 262, 280, 829, 997, 999, 1020, 1021, 1241, 1250

OCTETS • 146, 174, 265, 266, 282, 284, 839, 840, 1073, 1081, 1089, 1095, 1106, 1113, 1247

OF • 54, 55, 101, 133, 134, 147, 444, 446, 449, 450, 573, 580, 582, 625, 666, 667, 668, 669, 670, 671, 672, 824, 853, 866, 886, 917, 923, 930, 1044, 1046, 1048, 1052, 1105, 1106, 1192, 1213

OLD • 88, 147, 666, 667, 668, 716, 720

old row • 135, 941

old transition table • 133, 135

<old transition table name> • 135, **666**, 667, 669

old transition variable • 134

<old transition variable name> • 136, 666, **667**, 669

ON • 57, 69, 70, 75, 89, 147, 335, 539, 573, 579, 581, 582, 597, 598, 599, 601, 617, 623, 633, 647, 648, 651, 655, 659, 666, 671, 672, 673, 677, 680, 698, 714, 717, 720, 744, 769, 772, 773, 774, 777, 779, 878, 897, 955, 1135

one-to-one • 372

one-to-one with respect to • 373, 855, 886, 1048, 1052

ONLY • 17, 59, 76, 87, 101, 147, 247, 325, 333, 334, 377, 449, 628, 631, 687, 749, 750, 805, 807, 808, 853, 854, 855, 866, 867, 869, 872, 873, 879, 880, 881, 882, 886, 888, 889, 890, 904, 909, 911, 916, 918, 928, 930, 948, 949, 970, 1013, 1044, 1046, 1048, 1052, 1056, 1163, 1194

<only spec> • **325**, 330, 332, 334, 792, 793, 796, 799, 1194

OPEN • 147, 856, 1039

<open statement> • 101, 106, 110, 113, 115, 117, 807, 832, 849, **856**, 1060, 1061, 1064, 1065, 1132, 1251, 1264

operand of a grouping operation • 481

operand of an equality operation • 478

operand of an ordering operation • 485

operative data type correspondences • 489, 529, 731, 741

operative data type correspondences table • 529, 731

OPTION • 57, 60, 120, 121, 122, 146, 246, 332, 582, 625, 626, 627, 631, 632, 633, 668, 672, 673, 688, 771, 772, 773, 774, 775, 776, 777, 778, 781, 783, 786, 789, 790, 791, 792, 793, 795, 796, 799, 800, 802, 803, 869, 877, 880, 881, 882, 884, 888, 889, 890, 911, 924, 931, 1137, 1165, 1181, 1193, 1209

OPTIONS • 146, 574, 625

OR • 33, 73, 147, 213, 299, 300, 302, 308, 309, 404, 407, 411, 440, 883

ORDER • 147, 356, 377, 485, 548, 555, 556, 687, 749, 752, 1200, 1249

<order by clause> • 23, 29, 54, 63, 101, 198, 207, 208, 230, 312, 313, 329, 359, 376, **377**, 385, 390, 391, 634, 854, 855, 886, 1046, 1053, 1155, 1184, 1185, 1186, 1246, 1247, 1248, 1279, 1280

order of execution • 134, 932, 939

order-equivalent • 359, 360, 1248

ordered cursor • 101, 1053

<ordered set function> • 205, 209, 547, **548**

ORDERING • 146, 687, 744, 749, 753

<ordering category> • **749**

<ordering form> • **749**

ordering operation • 411, 485, 486, 751, 753, 754, 1218

<ordering specification> • 62, 359, 361, 362, 363, 556, **559**

ordering window • 62

ORDINALITY • 76, 146, 326, 327, 328, 329

ordinary column reference • 202

<ordinary grouping set> • **345**, 347, 348, 349, 350, 351, 353, 1217

<ordinary grouping set list> • **345**, 347, 348

original method • 41, 501, 503, 682, 686, 691, 697, 703, 706

<original method specification> • 40, 41, 42, **675**, 677, 680, 682, 685, 688, 697, 1188

originally-defined attributes • 677

originally-defined column • 578, 630

originally-defined columns • 59, 629

OTHERS • 146, 357, 365

OUT • 147, 715, 723, 726, 727, 729, 731, 1012



OUTER • 147, 335, 1134, 1135  
 <outer join type> • **335**, 336, 338, 1268  
 outer reference • 202, 207, 323, 336, 344, 354, 550, 557, 558, 1174, 1175, 1214  
 OUTPUT • 146, 1017  
 <output using clause> • **1027**, 1028, 1031, 1032, 1041, 1161, 1241, 1250  
 OVER • 147, 207, 208, 209, 555, 556  
 OVERLAPS • 39, 147, 439, 440, 1166, 1167  
 <overlaps predicate> • 39, 251, 401, 402, **439**, 440, 485, 1009, 1167  
 <overlaps predicate part 1> • 211, 212, **439**  
 <overlaps predicate part 2> • 212, **439**  
 OVERLAY • 21, 147, 274, 275, 287, 1005, 1215  
 <override clause> • **873**, 874, 875, 877, 878, 879, 1192  
 OVERRIDING • 146, 675, 679, 873, 875, 877, 879, 884  
 overriding method • 41, 89, 501, 523, 683, 684, 703, 704, 705  
 <overriding method specification> • 40, 41, 42, **675**, 682, 686, 703  
 owned by • 81, 121, 771, 772, 774

## — P —

P • 143, 146, 151, 176, 1156, 1207, 1283  
 PAD • 19, 25, 146, 408, 540, 652  
 <pad characteristic> • **652**, 653  
 PARAMETER • 148, 372, 527, 529, 531, 532, 534, 680, 698, 716, 722, 725, 730, 734, 736, 741, 742  
 <parameter mode> • 679, 698, 704, **715**, 723, 726, 727, 728, 729, 730, 731, 1012, 1022, 1138, 1139  
 parameter passing style • 90, 734, 741, 742  
 <parameter style> • 40, 42, 92, 680, 685, 686, 688, 701, 706, 716, **717**, 722, 736, 1188  
 <parameter style clause> • 676, 680, 698, **716**, 717, 720, 740, 741, 742  
 <parameter type> • 471, 679, 699, 704, 705, **715**, 722, 723, 726, 727, 728, 729, 730, 731, 737, 738, 1195, 1196, 1197, 1198, 1205  
 <parameter using clause> • 1023, **1032**, 1033  
 PARAMETER\_MODE • 146, 992, 997, 999, 1022, 1126, 1138, 1139, 1241  
 PARAMETER\_NAME • 146, 1126, 1138, 1139  
 PARAMETER\_ORDINAL\_POSITION • 146, 997, 999, 1019, 1022, 1126, 1138, 1139, 1241  
 PARAMETER\_SPECIFIC\_CATALOG • 146, 997, 999, 1019, 1022, 1241  
 PARAMETER\_SPECIFIC\_NAME • 146, 997, 999, 1019, 1022, 1241  
 PARAMETER\_SPECIFIC\_SCHEMA • 146, 997, 999, 1019, 1022, 1241  
 parametric column • 61

<parenthesized boolean value expression> • **299**, 300, 301  
 <parenthesized joined table> • 325, **326**  
 <parenthesized value expression> • **188**  
 part of • xx, xx, xx, xx, xxi, xxi, xxi, xxi, xxi, xxi, xxi, xxi, 1, 8, 9, 22, 24, 25, 26, 34, 35, 40, 47, 50, 54, 58, 62, 69, 80, 81, 98, 104, 126, 127, 129, 159, 358, 416, 418, 955, 957, 975, 1088, 1153, 1154, 1225, 1245, 1253, 1255, 1257, 1287  
 PARTIAL • 68, 69, 146, 436, 437, 597, 599, 933, 935, 936  
 partial identifier chain • 198, 367  
 <partial method specification> • 675, **676**, 689, 697, 703, 1188  
 participating datetime fields • 511  
 PARTITION • 148, 335, 356  
 partition dependency graph • 387  
 <partitioned join column reference> • **335**, 336  
 <partitioned join column reference list> • **335**  
 <partitioned join table> • **335**, 336, 338, 339, 340, 343, 1172  
 partitioning column • 357  
 PASCAL • 146, 489, 513, 529, 731, 739, 741, 809, 824, 825, 827, 1059, 1109, 1164, 1165  
 <Pascal array locator variable> • 1105, **1106**, 1108, 1109, 1197  
 <Pascal BINARY variable> • **1105**, 1107, 1110, 1203  
 <Pascal BLOB locator variable> • 1105, **1106**, 1108, 1110, 1206  
 <Pascal BLOB variable> • **1105**, 1107, 1110, 1206  
 <Pascal CLOB locator variable> • **1105**, 1108, 1110, 1206  
 <Pascal CLOB variable> • **1105**, 1106, 1107, 1110, 1206  
 <Pascal derived type specification> • **1105**  
 <Pascal host identifier> • 1057, **1105**, 1106, 1107, 1109  
 <Pascal multiset locator variable> • 1105, **1106**, 1108, 1110, 1198  
 <Pascal REF variable> • 1105, **1106**, 1109, 1191  
 <Pascal type specification> • **1105**, 1106, 1107  
 <Pascal user-defined type locator variable> • 1105, **1106**, 1108, 1110, 1196  
 <Pascal user-defined type variable> • 1105, **1106**, 1107, 1109, 1200  
 <Pascal variable definition> • 1056, **1105**, 1106, 1109  
 PATH • 146, 515, 978  
 <path column> • **393**, 394  
 <path specification> • **515**, 567, 736, 805, 1056, 1192  
 <path-resolved user-defined type name> • 84, 95, 130, 165, 173, **175**, 179, 184, 190, 192, 194, 233, 237, 239, 449, 516, 519, 520, 573, 577, 580, 581, 582, 625, 627, 632, 674, 717, 722, 726, 727, 729, 730, 744, 806, 902, 979, 1058, 1072, 1075, 1077, 1079, 1080, 1083, 1085, 1088, 1091, 1092, 1095, 1097, 1099, 1101, 1103, 1104,

1106, 1108, 1110, 1112, 1114, 1116, 1187, 1192, 1194, 1196, 1197, 1201

peers • 61, 63, 360, 361, 363, 364, 365, 560

<percent> • 22, 23, 139, **140**, 416, 417, 418, 420, 422, 423

PERCENT\_RANK • 63, 64, 66, 148, 207, 208, 210, 1221, 1227

PERCENTILE\_CONT • 66, 148, 548, 556, 1236

PERCENTILE\_DISC • 66, 148, 548, 556

<period> • 81, 139, **140**, 153, 155, 157, 161, 162, 163, 198, 200, 202, 221, 223, 232, 235, 274, 366, 368, 511, 516, 676, 821, 892, 1087

permitted • 30, 43, 47, 55, 64, 66, 91, 101, 102, 122, 123, 130, 133, 320, 1146, 1147, 1228, 1241

persistent • **6**

<PL/I array locator variable> • 1111, **1112**, 1115, 1116, 1197

<PL/I BINARY variable> • **1111**, 1113, 1116, 1203

<PL/I BLOB locator variable> • 1111, **1112**, 1114, 1116, 1206

<PL/I BLOB variable> • **1111**, 1113, 1116, 1206

<PL/I CLOB locator variable> • 1111, **1112**, 1114, 1116, 1206

<PL/I CLOB variable> • **1111**, 1113, 1116, 1206

<PL/I derived type specification> • **1111**

<PL/I host identifier> • 1057, **1111**, 1112, 1113, 1114, 1115

<PL/I multiset locator variable> • 1111, **1112**, 1115, 1116, 1198

<PL/I REF variable> • 1111, **1112**, 1115, 1116, 1191

<PL/I type fixed binary> • 1111, **1112**

<PL/I type fixed decimal> • 1111, **1112**

<PL/I type float binary> • 1111, **1112**

<PL/I type specification> • **1111**, 1112, 1113, 1242

<PL/I user-defined type locator variable> • 1111, **1112**, 1114, 1116, 1197

<PL/I user-defined type variable> • 1111, **1112**, 1114, 1116, 1200

<PL/I VARBINARY variable> • **1111**, 1114, 1117, 1203

<PL/I variable definition> • 1056, **1111**, 1112, 1115, 1242

PLACING • 146, 274, 275, 1005

PLI • 146, 489, 513, 529, 731, 739, 741, 809, 824, 825, 826, 827, 828, 829, 1059, 1164, 1165

<plus sign> • 23, 139, **140**, 144, 149, 153, 254, 255, 288, 290, 293, 420, 422, 423, 1006, 1166

Popping • 97

POSITION • 8, 148, 257, 1005

<position expression> • 22, 23, 29, 31, **256**, 258, 478, 479, 1005, 1232, 1260, 1283

POSITION\_REGEX • 23, 148, 256, 1255

positioned delete statement • 103

positioned update statement • 103

possible scope tags • 198

possibly candidate routine • 516, 517

possibly contains SQL • 40, 42, 91, 92, 132, 525, 527, 530, 686, 701, 732, 736

possibly modifies SQL-data • 40, 42, 91, 212, 385, 525, 527, 530, 531, 602, 661, 668, 686, 701, 725, 732, 734, 736, 868, 880, 888, 941, 1237, 1238

possibly modify SQL-data • 110, 746, 750, 756, 759

possibly non-deterministic • 62, 79, 91, 92, 212, 251, 252, 301, 320, 331, 332, 337, 338, 370, 384, 627, 725, 732, 833, 863, 1238

possibly nullable • 52, 338, 339, 371, 383, 621, 645, 664, 1019

possibly read SQL-data • 104, 110, 746

possibly reads SQL-data • 40, 91, 92, 132, 525, 527, 530, 531, 585, 686, 701, 732, 734

potentially recursive • 377, 378, 387

potentially referenced cursors • 1050

potentially updatable • 373

potentially updatable column • 373

POWER • 148, 258, 260, 554

<power function> • 32, 256, **258**, 260, 263, 266, 1222, 1232

precede • 63, 360, 361, 364, 511, 512, 560, 807, 825, 1058

PRECEDING • 146, 208, 357, 358, 361, 362, 363, 364

PRECISION • 14, 15, 30, 148, 174, 178, 179, 183, 469, 474, 840, 842, 843, 844, 846, 847, 984, 985, 986, 997, 1000, 1001, 1002, 1021, 1076, 1085, 1094, 1099, 1230, 1231, 1250

<precision> • 30, 85, **174**, 178, 822, 838, 1004, 1100, 1101, 1111, 1116, 1230, 1231

predefined data types • 13, 14, 16, 979, 1025, 1029

<predefined type> • **173**, 638, 674, 675, 676, 677, 1072, 1079, 1088, 1095, 1101, 1106, 1112, 1273

<predicate> • 33, 299, 300, 301, **401**, 402, 416, 418, 426, 1262

preferred candidate key • 80

<preparable dynamic cursor name> • 1011, 1048, **1050**, 1051, 1052

<preparable dynamic delete statement: positioned> • 102, 103, 107, 111, 114, 115, 117, 870, 872, 882, 890, 909, 911, 938, 1003, **1048**, 1049, 1050, 1132, 1162

<preparable dynamic update statement: positioned> • 102, 103, 107, 111, 114, 115, 117, 870, 872, 882, 890, 909, 911, 938, 1003, 1050, **1052**, 1053, 1132, 1162

<preparable implementation-defined statement> • 115, 1003, **1004**, 1241

<preparable SQL control statement> • **1003**, 1004

<preparable SQL data statement> • **1003**

<preparable SQL schema statement> • **1003**

<preparable SQL session statement> • 1003, **1004**

<preparable SQL transaction statement> • **1003**

<preparable statement> • 86, 91, 95, 129, 130, 133, 164, 165, 166, 167, 192, 487, 518, 519, 733, 853, 974, 975, 977, 978, **1003**, 1004, 1011, 1228, 1229

PREPARE • 148, 531, 1003, 1012, 1015, 1034

<prepare statement> • 10, 11, 83, 86, 87, 95, 109, 111, 129, 130, 164, 165, 166, 167, 170, 192, 487, 518, 519, 833, 853, 855, 902, 974, 975, 977, 978, 979, **1003**, 1012, 1013, 1015, 1017, 1032, 1033, 1035, 1050, 1132, 1161, 1229, 1241, 1255

prepared • xix, 10, 83, 85, 86, 87, 91, 95, 100, 101, 102, 104, 111, 118, 123, 124, 129, 130, 131, 165, 166, 167, 170, 192, 194, 487, 531, 808, 853, 855, 901, 902, 956, 957, 958, 974, 975, 977, 978, 979, 991, 1004, 1012, 1015, 1017, 1018, 1021, 1023, 1027, 1032, 1033, 1034, 1036, 1037, 1039, 1129, 1130, 1227, 1228, 1229, 1235, 1240, 1242, 1251, 1255

*prepared statement not a cursor specification* • 1037, 1039, 1147, 1255

PRESERVE • 57, 146, 573, 581, 601

PRIMARY • 52, 68, 73, 148, 578, 595, 596, 598, 618

<primary datetime field> • 34, 35, 38, 39, 158, 159, 179, 180, 181, 216, 224, 225, 226, 227, 228, 257, 259, 260, 261, 288, 289, 290, 295, 296, 361, 409, 439, 464, **509**, 510, 511, 1232

primary effect • 59, 60, 111

primary key • 57, 68, 73, 80, 581, 1018, 1020

primitive • 47, 346, 347, 352

PRIOR • 146, 857, 905, 906

<privilege column list> • 776, **779**, 780, 781, 786, 787, 1181, 1214, 1262, 1263

privilege dependency graph • 789

privilege descriptor • 81, 120, 121, 539, 561, 562, 563, 564, 565, 566, 582, 605, 632, 633, 639, 648, 649, 653, 657, 670, 671, 672, 673, 688, 702, 707, 714, 733, 767, 772, 774, 777, 781, 787, 788, 789, 790, 791, 800, 803, 1166

<privilege method list> • 120, 121, 776, **779**, 780, 781, 786, 787, 1189

*privilege not granted* • 776, 777, 1151

*privilege not revoked* • 802, 1151

PRIVILEGES • 146, 623, 777, 779, 780, 802

<privileges> • 776, 777, **779**, 780, 781, 786, 787, 802, 803, 1165, 1166, 1187, 1189, 1193, 1263

PROCEDURE • 148, 538, 541, 542, 543, 715, 720, 811, 1036

<procedure name> • **162**, 811, 812, 814, 1060, 1061, 1251

*prohibited SQL-statement attempted* • 526, 527, 530, 1150

*prohibited SQL-statement attempted* • 1148

*prohibited statement encountered during trigger execution* • 941, 1149

proper subtables • 58, 59

proper subtype • 9, 43, 44, 47, 468, 682, 684, 694, 704, 710, 756

proper subview • 627

proper supertype • 9, 40, 41, 43, 682, 683, 703, 704, 705

proper superview • 627

property • 6

PUBLIC • 118, 122, 146, 168, 539, 648, 779, 780, 788, 789, 790, 898, 972

Pushing • 97

## — Q —

<qualified asterisk> • **366**, 369, 371, 1261

<qualified identifier> • 89, 161, **162**, 163, 164, 165, 166, 167, 199, 201, 202, 239, 241, 286, 330, 367, 374, 449, 516, 517, 522, 529, 676, 677, 678, 679, 682, 684, 693, 697, 700, 703, 705, 709, 710, 724, 725, 897, 1136, 1137, 1215

<qualified join> • **335**, 336, 379

qualifying query • 79, 202, 203, 205, 1246

<quantified comparison predicate> • 23, 29, 251, 301, 372, 401, 402, **430**, 431, 478, 485, 1009, 1262

<quantified comparison predicate part 2> • 211, **430**

<quantifier> • 301, **430**, 431

queried column reference • 202

<query expression> • vi, 54, 56, 57, 58, 59, 63, 66, 71, 75, 78, 79, 90, 110, 112, 198, 205, 207, 208, 230, 252, 301, 306, 307, 312, 329, 330, 331, 333, 336, 338, 344, 354, 355, 358, 359, 369, 370, 372, 373, 374, **376**, 377, 378, 379, 380, 381, 384, 385, 387, 388, 389, 391, 392, 393, 398, 399, 463, 550, 551, 561, 563, 565, 582, 585, 601, 602, 605, 610, 611, 612, 616, 622, 625, 626, 627, 630, 631, 633, 634, 635, 646, 650, 654, 659, 661, 668, 669, 671, 695, 712, 713, 740, 743, 747, 748, 753, 754, 772, 774, 788, 789, 791, 794, 795, 797, 798, 799, 833, 853, 856, 870, 873, 875, 877, 884, 885, 891, 896, 902, 918, 920, 921, 923, 930, 1008, 1039, 1048, 1052, 1123, 1155, 1169, 1178, 1181, 1183, 1184, 1185, 1186, 1189, 1210, 1214, 1248, 1249, 1267, 1269, 1270, 1279, 1280

<query expression body> • 58, 169, 329, 359, **376**, 377, 378, 379, 380, 381, 382, 383, 384, 385, 388, 393, 481, 915, 921, 927, 1248

*query expression too long for information schema* • 633, 1151

<query name> • 54, 56, 58, 79, **162**, 169, 170, 326, 330, 331, 332, 334, 376, 377, 378, 379, 380, 381, 387, 388, 393, 564, 1210

query name dependency graph • 377

query name in scope • 377

<query primary> • 79, **376**, 379, 381, 382, 383, 384, 385, 391, 915, 921, 927, 1186

<query specification> • vi, 23, 29, 54, 59, 62, 63, 66, 78, 101, 202, 205, 207, 209, 230, 231, 252, 329, 338, 346, 359, **366**, 367, 369, 370, 371, 372, 373, 374, 376, 379, 380, 381, 384, 385, 386, 388, 393, 432, 481, 563, 619,

627, 628, 663, 833, 854, 863, 915, 916, 921, 928, 1042, 1184, 1219, 1248, 1270  
 <query term> • 79, **376**, 379, 380, 381, 382, 383, 385, 388, 391, 481, 915, 927, 1169, 1248  
 <question mark> • 23, 85, 140, **141**, 191, 420, 422, 423  
 <quote> • 139, **140**, 144, 145, 149, 150, 152, 153, 154, 155, 156, 157, 160, 1169, 1204, 1259  
 <quote symbol> • 150, **152**, 157

## — R —

RANGE • 61, 148, 208, 356, 358, 361  
 range variable • 59, 134, 136, 199, 329, 330, 333, 372, 580  
 RANK • 63, 64, 66, 148, 207, 208, 550, 1236  
 <rank function type> • **207**, 372, 548, 555  
 READ • 101, 124, 125, 127, 146, 853, 854, 855, 904, 948, 949, 970  
 read-only • 124, 127, 835, 836, 869, 871, 876, 881, 889, 908, 910, 948, 970  
*read-only SQL-transaction* • 869, 871, 876, 881, 889, 908, 910, 1121, 1149  
 read-write • 124  
 readily known not null • 338, 371  
 readily-known-not-null condition • 52, 300  
 reading SQL-data • 130  
*reading SQL-data not permitted* • 525, 527, 530, 1148, 1150  
 READS • 148, 717, 725, 732, 736, 742, 1237  
 REAL • 14, 15, 30, 148, 174, 178, 179, 183, 469, 474, 815, 823, 825, 828, 840, 841, 843, 844, 845, 847, 984, 986, 1076, 1085, 1094, 1098, 1100, 1101, 1105, 1109, 1230, 1231  
 received cursor • 9, 93, 99, 100, 101, 102, 526, 901, 904, 907, 1037, 1228, 1239  
 RECURSIVE • 79, 148, 170, 327, 334, 376, 377, 391, 625, 626, 627, 633, 1209, 1210  
 recursive • 300, 377, 378, 380, 383, 387, 388, 393, 494, 626, 668  
 recursive query names in scope • 388  
 <recursive search order> • **393**, 394, 1255  
 recursively referred to • 383, 384  
 redundant duplicates • **9**, 443  
 REF • 9, 13, 46, 47, 58, 148, 175, 183, 470, 573, 577, 609, 610, 627, 628, 629, 630, 675, 687, 688, 714, 840, 842, 843, 845, 846, 847, 876, 883, 984, 986, 1084  
 REF value • **9**, 46  
 <reference column list> • **597**, 598, 616  
 <reference generation> • 573, **574**, 582, 1191  
 <reference resolution> • 46, 188, 189, **246**, 247, 372, 610, 622, 632, 635, 771, 772, 773, 792, 793, 795, 796, 799, 1191

<reference type> • 7, 9, 16, 17, 42, 44, 46, 47, 173, **175**, 179, 180, 183, 184, 216, 217, 218, 219, 229, 233, 234, 241, 243, 244, 246, 250, 442, 454, 457, 458, 465, 470, 577, 579, 589, 609, 610, 611, 619, 622, 623, 627, 628, 629, 630, 635, 636, 691, 693, 695, 712, 714, 721, 745, 746, 799, 800, 876, 884, 1021, 1072, 1075, 1080, 1084, 1088, 1091, 1095, 1098, 1101, 1103, 1104, 1106, 1109, 1112, 1115, 1189, 1190, 1191, 1194, 1195, 1201, 1226  
 <reference type specification> • 674, **675**, 676, 677, 689, 1190, 1192  
 <reference value expression> • 243, 246, **250**, 252, 253, 610, 622, 635, 1190  
 reference-ordered • 17, 485  
 referenceable table • 46, 47, 56, 58, 243, 465, 579, 580, 604, 609, 616, 619, 623, 628, 629, 792  
 referenceable view • 627  
 <referenceable view specification> • **625**, 626, 627, 629, 633, 1191  
 referenced column • 68, 69, 598  
 referenced columns • 68, 69, 598, 599, 600, 619  
 referenced table • 59, 68, 69, 70, 133, 487, 597, 598, 599, 619, 622, 871, 933, 934, 935, 937  
 <referenced table and columns> • 68, **597**, 598, 599, 600, 616  
 <referenced type> • **9**, 16, 46, 47, **175**, 179, 183, 241, 243, 404, 454, 458, 470, 609, 619, 629, 691, 693, 695, 712, 714, 1075, 1084, 1091, 1098, 1104, 1109, 1115, 1189  
 referenced value • **9**  
 REFERENCES • 120, 148, 203, 332, 581, 582, 597, 599, 605, 616, 617, 632, 639, 771, 773, 774, 777, 779, 780, 781, 787, 788, 789, 791, 792, 793, 794, 796, 898  
 <references specification> • 584, 586, **597**, 600, 1181, 1266  
 REFERENCING • 148, 666, 667  
 referencing column • 69, 598, 934  
 <referencing columns> • 68, 69, 70, **597**, 598, 599, 600, 934, 935, 936, 937, 1212  
 referencing table • 69, 70, 597, 598, 599, 933, 935  
 referencing type • 41, 42, 46, 678, 685  
 referent • 90, 198, 199, 200, 201, 202, 203, 367, 374, 1215  
 <referential action> • **597**, 599, 600, 1212  
 referential actions • 68, 69, 125, 836, 933  
 <referential constraint definition> • 23, 29, 69, 481, 487, 593, **597**, 599, 1266, 1272  
 referential delete action • 69, 70, 934, 935, 1266  
 <referential triggered action> • 68, **597**, 600  
 referential update action • 69, 70, 936, 937, 1266  
 refinable • 198, 199, 200, 367  
 <regex capture group> • 256, **257**, 259, 265, 273, 275, 276, 281, 1011  
 <regex like predicate> • 23, 401, 402, **426**, 427, 1011, 1184, 1185

- <regex like predicate part 2> • 211, **426**
- <regex occurrence> • 256, **257**, 259, 265, 273, 274, 275, 276, 281, 284, 1011
- <regex occurrences function> • **256**, 258, 265, 267, 1011, 1184, 1185, 1232, 1247
- <regex position expression> • 23, **256**, 258, 259, 265, 266, 267, 1011, 1184, 1185, 1232, 1247
- <regex position start or after> • 256, **257**, 259, 266
- <regex subject string> • **256**, 258, 259, 265, 273, 274, 276, 277, 281, 284, 1011
- <regex substring function> • 21, **273**, 275, 276, 281, 282, 287, 1011, 1185, 1247
- <regex transliteration> • 21, 273, **274**, 275, 277, 281, 284, 285, 287, 1011, 1185, 1233, 1247
- <regex transliteration occurrence> • **274**, 277, 284, 285
- REGR\_AVGX • 66, 148, 548, 553
- REGR\_AVGY • 66, 148, 548, 553
- REGR\_COUNT • 65, 148, 548, 550, 553, 1236
- REGR\_INTERCEPT • 65, 148, 548, 554
- REGR\_R2 • 148, 548
- REGR\_SLOPE • 65, 148, 548, 554
- REGR\_SXX • 66, 148, 548, 553
- REGR\_SXY • 66, 148, 548, 553
- REGR\_SYY • 66, 148, 548, 553
- <regular character set> • 420, **421**
- <regular character set identifier> • **421**, 422
- <regular expression> • 282, **420**, 422, 423
- <regular expression substring function> • 20, 21, **273**, 275, 276, 281, 282, 283, 287, 1220
- <regular factor> • **420**
- <regular identifier> • 24, **143**, 149, 150, 151, 161, 1171, 1225, 1229, 1260
- <regular primary> • **420**
- <regular term> • **420**
- <regular view specification> • **625**, 630
- RELATIVE • 41, 45, 146, 404, 406, 442, 749, 751, 857, 905, 906, 1234
- <relative category> • **749**, 750
- <relative function specification> • **749**, 750
- RELEASE • 148, 954
- <release savepoint statement> • 89, 108, 123, 527, 530, 725, 832, **954**, 1132, 1213, 1238
- released • 85
- <repeat argument> • 59, **325**, 326, 333, 1227
- <repeat factor> • **420**, 423
- REPEATABLE • 124, 125, 127, 146, 325, 948, 949
- <repeatable clause> • 59, **325**, 333, 1227
- repertoire • **5**
- replacement set • 882, 883, 890, 911, 912, 925, 930, 934, 935, 936, 937
- <representation> • **674**, 676, 1273
- representation type of the referencing type • 677
- represented • 5, 18, 31, 85, 120, 130, 148, 149, 155, 156, 157, 159, 193, 255, 387, 405, 417, 423, 592, 669, 899, 986, 987, 997, 1024, 1225
- request failed* • 869, 871, 876, 881, 889, 908, 910, 1144
- request rejected* • 903, 1144
- <reserved word> • 146, **147**, 150, 164, 1255, 1258, 1259, 1282
- respective values • 226, 441
- RESTART • 146, 768
- RESTRICT • 70, 146, 570, 597, 600, 610, 616, 619, 622, 623, 635, 646, 654, 663, 695, 708, 712, 740, 743, 747, 753, 761, 763, 769, 800, 934, 935, 936, 937, 1212
- restrict violation* • 934, 935, 936, 937, 1148
- restricted data type attribute violation* • 1025, 1026, 1028, 1029, 1030, 1147, 1241
- RESULT • 9, 42, 148, 675, 677, 679, 680, 681, 688, 697, 699, 700, 704, 715, 716, 720, 722, 723, 737, 1188
- <result> • 211, **212**, 213, 214
- <result cast> • 93, 532, 533, 680, 681, 682, 698, 699, 703, **716**, 718, 719, 720, 722, 726, 734, 736
- <result cast from type> • 40, 42, 681, 685, 686, 688, 689, 698, 701, 706, **716**, 721, 722, 726, 1219, 1220
- result data items • 489, 490, 491, 528
- result data type • **9**
- <result expression> • **212**, 213, 1007, 1008
- result set • 9, 10, 57, 93, 94, 100, 101, 102, 103, 132, 526, 536, 538, 853, 857, 901, 903, 904, 905, 907, 909, 1017, 1036, 1037, 1047, 1048, 1052, 1053, 1246
- <result set cursor> • 99, 101, **1036**, 1037, 1038, 1218
- result set descriptor • 94
- result set sequence • **9**, 93, 101, 904
- result sets returned* • 537, 1151
- result SQL parameter • **9**
- <result using clause> • 1027, **1032**, 1033, 1162
- retrospectively deterministic • 301
- RETURN • 10, 94, 100, 102, 148, 686, 687, 688, 751, 851, 902, 903, 907, **944**, 1038
- <return statement> • 108, 112, 114, 116, 117, 527, 832, **944**, 1132, 1273
- <return value> • **944**, 1026, 1030
- returned result set • **9**, 93
- <returned result sets characteristic> • 94, **716**, 720, 737, 740, 741, 1218
- returned value • 29, 85, 527, 944
- RETURNED\_CARDINALITY • 146, 997, 999, 1031
- RETURNED\_LENGTH • 146, 997, 999, 1031
- RETURNED\_OCTET\_LENGTH • 146, 997, 999, 1031
- RETURNED\_SQLSTATE • 146, 537, 1126, 1136, 1137, 1138, 1139

RETURNS • 89, 148, 677, 686, 687, 688, 716, 717, 751  
 <returns clause> • 676, 680, 682, 688, 689, 698, 701, 703, 706, 715, **716**, 719, 720, 944, 1219, 1220  
 <returns data type> • 40, 42, 92, 479, 485, 532, 533, 650, 654, 679, 680, 681, 683, 684, 685, 686, 688, 689, 697, 698, 699, 701, 704, 705, 706, **716**, 719, 721, 722, 723, 726, 727, 729, 730, 731, 734, 735, 737, 738, 944, 1030, 1195, 1196, 1197, 1198, 1205, 1219, 1220  
 <returns table type> • **716**, 718, 738, 1216  
 <returns type> • 657, **716**, 718  
 <reverse solidus> • **140**, 149, 1229  
 REVOKE • 148, 617, 623, 647, 651, 655, 659, 671, 672, 673, 714, 744, 769, 785, **786**, 802, 803, 1165, 1166  
 revoke destruction action • 791, 796, 797  
 <revoke option extension> • **786**, 802, 803, 1165, 1193  
 <revoke privilege statement> • **786**, 787, 791, 800, 802, 1132  
 <revoke role statement> • 785, **786**, 787, 791, 800, 802, 1132, 1216  
 <revoke statement> • 105, 616, 617, 623, 635, 636, 647, 651, 655, 659, 664, 671, 672, 673, 713, 714, 743, 744, 769, 780, **786**, 787, 790, 791, 800, 802, 803, 832, 1165, 1166, 1267  
 RIGHT • 74, 75, 76, 148, 335, 336, 338, 339, 341, 379, 380, 1134, 1135  
 <right arrow> • **145**, 241  
 <right brace> • 140, **141**, 420  
 <right bracket> • 23, 140, **141**, 421, 422, 424, 1078, 1105  
 <right bracket or trigraph> • **141**, 175, 191, 196, 248, 306, 312, 892  
 <right bracket trigraph> • **141**, 145  
 <right paren> • 23, 139, **140**, 173, 174, 175, 188, 191, 200, 205, 207, 211, 215, 233, 235, 237, 246, 249, 256, 257, 258, 273, 274, 275, 291, 293, 298, 299, 306, 311, 315, 325, 326, 335, 345, 356, 366, 376, 381, 398, 412, 420, 422, 424, 449, 509, 516, 541, 547, 548, 573, 574, 584, 595, 597, 601, 625, 631, 661, 666, 674, 675, 715, 716, 745, 747, 755, 758, 759, 761, 779, 811, 866, 873, 878, 892, 1004, 1005, 1055, 1057, 1071, 1072, 1079, 1080, 1087, 1088, 1094, 1095, 1100, 1101, 1105, 1111, 1112, 1261, 1263  
 <rights clause> • **717**, 725, 738, 1215  
 ROLE • 146, 782, 785, 972  
 role authorization • 122, 781, 782, 783, 787, 790, 791, 800, 804, 972  
 role authorization descriptor • 122  
 <role definition> • 106, 121, 568, **782**, 831, 1132, 1216, 1217, 1238  
 role dependency graph • 790  
 <role granted> • **783**, 784  
 <role name> • 121, 161, **163**, 169, 170, 569, 782, 783, 784, 785, 786, 788, 790, 972, 1216

<role revoked> • **786**, 787  
 <role specification> • **972**, 1010  
 ROLLBACK • 125, 126, 148, 957  
 <rollback statement> • 68, 89, 97, 101, 108, 110, 112, 123, 124, 126, 127, 527, 530, 725, 808, 832, **957**, 958, 959, 966, 1120, 1132, 1213, 1228, 1238, 1251, 1266  
 ROLLUP • 63, 73, 78, 148, 345, 347  
 <rollup list> • **345**, 346, 347, 349, 353, 1217  
 root data type • 180  
 ROUTINE • 146, 541, 542, 543, 571, 611, 620, 624, 637, 664, 713, 744, 748, 753, 762, 764, 765, 802  
 routine authorization identifier • 91, 92, 526, 735, 736  
 <routine body> • 88, 89, 90, 91, 715, **717**, 723, 733, 774, 797, 944  
 <routine characteristic> • **716**, 738, 739, 1164, 1165  
 <routine characteristics> • 715, **716**, 720, 738, 1214  
 routine descriptor • 81, 91, 518, 531, 532, 534, 571, 610, 612, 616, 620, 622, 624, 635, 636, 637, 646, 650, 654, 657, 659, 663, 664, 679, 682, 695, 698, 703, 712, 713, 721, 733, 734, 735, 736, 737, 738, 740, 741, 742, 743, 747, 750, 753, 769, 797, 800, 1215, 1223  
 routine execution context • 131, 132, 518, 519, 525, 1037  
 <routine invocation> • 82, 84, 89, 90, 91, 94, 95, 130, 133, 188, 189, 212, 235, 236, 237, 239, 244, 252, 292, 326, 344, 372, 379, 385, 472, **516**, 535, 538, 557, 568, 585, 602, 631, 632, 661, 668, 695, 713, 725, 733, 735, 736, 738, 740, 743, 771, 772, 773, 774, 792, 793, 794, 796, 797, 806, 833, 868, 880, 888, **943**, 978, 1010, 1012, 1021, 1162, 1174, 1223, 1237, 1238, 1273, 1287  
 <routine name> • 84, 89, 90, 91, 94, 95, 130, 133, 198, 199, 201, 367, 374, 406, 407, **516**, 517, 522, 568, 700, 703, 705, 706, 709, 710, 724, 806, 978, 1063, 1138, 1139, 1215  
 routine name text item • 528, 529  
 routine SQL-path • 91, 92, 132, 518, 519, 526, 735, 736  
 <routine type> • **541**, 542, 543, 1188  
 ROUTINE\_CATALOG • 146, 1126, 1138, 1139  
 ROUTINE\_NAME • 146, 1126, 1138, 1139  
 ROUTINE\_SCHEMA • 146, 1126, 1138, 1139  
 ROW • 13, 148, 175, 208, 219, 312, 315, 316, 317, 357, 358, 361, 363, 364, 365, 377, 394, 395, 471, 491, 666, 667, 668, 670, 679, 699, 718, 723, 840, 842, 843, 845, 846, 847, 893, 984, 986, 998, 1000, 1021, 1024, 1025, 1030, 1208, 1212  
 row • **6**  
 <row subquery> • 59, 315, 316, 317, **398**, 399, 1178  
 <row type> • 9, 13, 16, 44, 46, 52, 53, 54, 55, 58, 94, 173, **175**, 180, 183, 184, 186, 198, 217, 219, 232, 316, 320, 322, 323, 328, 329, 331, 354, 367, 368, 372, 373, 388, 398, 430, 436, 442, 454, 457, 458, 462, 464, 465, 470, 471, 478, 479, 485, 498, 558, 582, 589, 603, 627, 857, 858, 860, 862, 983, 984, 986, 1008, 1009, 1021, 1207, 1208

<row type body> • **175**  
 <row value constructor> • 59, **315**, 316, 317, 318, 319, 1247  
 <row value constructor element> • 73, 300, **315**, 316, 317, 414, 1009, 1178  
 <row value constructor element list> • **315**, 1247  
 <row value constructor predicand> • 73, 214, 299, **315**, 316, 317, 318, 319, 414, 421, 426, 434, 444, 446, 448, 449, 1009, 1168, 1204  
 <row value expression> • 231, 250, 252, **318**, 319, 320, 412, 413, 478, 483, 491, 1008, 1177, 1208, 1272, 1273  
 <row value expression list> • **320**  
 <row value predicand> • 73, 211, 212, 213, 214, 251, 299, **318**, 319, 372, 403, 411, 412, 414, 419, 420, 421, 426, 428, 430, 434, 436, 439, 441, 443, 444, 446, 448, 449, 450, 478, 479, 485, 1009, 1168, 1169  
 <row value predicand 1> • **439**, 440, 1009  
 <row value predicand 2> • **439**, 440, 1009  
 <row value predicand 3> • **441**  
 <row value predicand 4> • **441**  
 <row value special case> • **318**, 319, 874, 1208  
 row-level trigger • 133, 134, 135, 136, 913, 920, 926, 940  
 ROW\_COUNT • 146, 1125, 1134, 1135, 1251  
 ROW\_NUMBER • 64, 148, 207, 208, 210, 252, 372, 556, 1221  
 ROWS • 57, 61, 148, 208, 252, 356, 359, 363, 377, 573, 579, 581, 598, 601, 897, 1011  
 rules of deduction • 72, 74

## — S —

S • 1088  
 S-ordered • 16  
 Feature S023, “Basic structured types” • 171, 184, 236, 240, 538, 688, 689, 691, 737, 781, 1186, 1187  
 Feature S024, “Enhanced structured types” • 238, 479, 482, 484, 486, 543, 688, 689, 691, 692, 737, 744, 778, 781, 877, 885, 896, 1187, 1188, 1189  
 Feature S025, “Final structured types” • 689, 1189  
 Feature S026, “Self-referencing structured types” • 691, 1189  
 Feature S027, “Create method by specific method name” • 738, 1190  
 Feature S028, “Permutable UDT options list” • 689, 1190  
 Feature S041, “Basic reference types” • 184, 242, 243, 253, 1077, 1085, 1092, 1099, 1104, 1109, 1116, 1190, 1191  
 Feature S043, “Enhanced reference types” • 184, 229, 245, 247, 582, 609, 611, 633, 689, 877, 1191, 1192  
 Feature S051, “Create table of type” • 582, 1192  
 Feature S071, “SQL paths in function and type name resolution” • 194, 515, 569, 592, 808, 978, 1066, 1192

Feature S081, “Subtables” • 582, 778, 781, 803, 1192, 1193  
 Feature S091, “Basic array support” • 184, 197, 248, 266, 305, 307, 333, 896, 1193, 1194, 1200, 1201  
 Feature S092, “Arrays of user-defined types” • 184, 1194  
 Feature S094, “Arrays of reference types” • 184, 1194  
 Feature S095, “Array constructors by query” • 307, 1194  
 Feature S096, “Optional array bounds” • 184, 1194  
 Feature S097, “Array element assignment” • 194, 1194  
 Feature S098, “ARRAY\_AGG” • 558, 1194  
 Feature S111, “ONLY in query expressions” • 334, 867, 1194  
 Feature S151, “Type predicate” • 450, 1194, 1195  
 Feature S161, “Subtype treatment” • 234, 1195  
 Feature S162, “Subtype treatment for references” • 234, 1195  
 Feature S201, “SQL-invoked routines on arrays” • 538, 737, 1195  
 Feature S202, “SQL-invoked routines on multisets” • 538, 737, 1195  
 Feature S211, “User-defined cast functions” • 746, 748, 1195, 1196  
 Feature S231, “Structured type locators” • 737, 738, 813, 1077, 1085, 1092, 1099, 1104, 1110, 1116, 1196, 1197  
 Feature S232, “Array locators” • 738, 813, 1077, 1085, 1092, 1099, 1104, 1109, 1116, 1197  
 Feature S233, “Multiset locators” • 738, 813, 1077, 1085, 1092, 1099, 1104, 1110, 1116, 1197, 1198  
 Feature S241, “Transform functions” • 194, 737, 757, 765, 808, 979, 1066, 1077, 1085, 1092, 1099, 1104, 1109, 1116, 1198, 1199, 1200  
 Feature S242, “Alter transform statement” • 758, 1200  
 Feature S251, “User-defined orderings” • 752, 754, 1200  
 Feature S261, “Specific type method” • 287, 1200  
 Feature S271, “Basic multiset support” • 184, 197, 249, 266, 311, 313, 333, 445, 448, 558, 1193, 1200, 1201  
 Feature S272, “Multisets of user-defined types” • 184, 1201  
 Feature S274, “Multisets of reference types” • 184, 1201  
 Feature S275, “Advanced multiset support” • 310, 447, 479, 558, 1201, 1202  
 Feature S281, “Nested collection types” • 184, 1202  
 Feature S291, “Unique constraint on entire row” • 596, 1202  
 Feature S301, “Enhanced UNNEST” • 334, 1202  
 <sample clause> • 59, **325**, 332, 333, 334, 1221, 1227  
 <sample method> • 59, **325**, 333  
 <sample percentage> • 59, **325**, 333  
 satisfied • 20, 59, 67, 68, 69, 71, 215, 216, 369, 542, 596, 733, 808, 833, 836, 937, 955, 1227, 1251  
 save area data item • 490, 491, 528  
 SAVEPOINT • 88, 148, 716, 720, 734, 953, 954, 957

- savepoint • **9**
- <savepoint clause> • 89, 97, 527, 530, 725, **957**, 958, 1214, 1238
- savepoint exception • 953, 954, 958, 1150
- <savepoint level indication> • **716**, 720, 738, 1214
- savepoint levels • 123, 734, 956, 957
- <savepoint name> • 123, **163**, 169, 170, 953, 954, 958, 1213
- <savepoint specifier> • 124, **953**, 954, 957, 1228
- <savepoint statement> • 89, 108, 110, 112, 123, 527, 530, 725, 832, **953**, 1132, 1213, 1238, 1240
- <scalar subquery> • 188, 189, 243, 249, **398**, 399, 412, 555, 556, 588, 594, 595, 599, 792, 793, 795, 796, 798, 1272
- SCALE • 146, 984, 985, 997, 1000, 1001, 1002, 1021, 1241, 1250
- <scale> • 30, 85, **174**, 178, 822, 838, 1100, 1101, 1111, 1116, 1230
- SCHEMA • 146, 567, 570, 805, 806, 975, 1056, 1060, 1241
- schema and data statement mixing not supported* • 834, 1122, 1149
- <schema authorization identifier> • 119, **567**, 568, 569, 1216
- <schema character set or path> • **567**
- <schema character set specification> • 186, **567**, 568, 569, 586, 801, 1175, 1237
- <schema definition> • 81, 95, 105, 119, 164, 165, 166, 167, 230, 246, 518, 519, 539, **567**, 568, 569, 575, 586, 626, 629, 638, 648, 649, 652, 656, 657, 661, 667, 676, 718, 721, 725, 732, 733, 767, 801, 831, 1132, 1237, 1267, 1272
- <schema element> • **567**, 569, 1272
- <schema function> • **715**
- <schema name> • 81, 82, 84, 91, 94, 95, 129, **161**, 162, 163, 164, 165, 166, 167, 168, 169, 189, 193, 235, 237, 239, 406, 515, 516, 518, 519, 521, 539, 540, 541, 567, 568, 569, 570, 571, 575, 578, 581, 586, 593, 603, 623, 626, 628, 629, 636, 638, 639, 641, 646, 648, 650, 652, 654, 656, 657, 659, 661, 664, 667, 668, 669, 671, 676, 679, 687, 688, 697, 698, 703, 708, 713, 718, 719, 720, 721, 732, 740, 741, 743, 744, 747, 748, 750, 751, 753, 767, 769, 805, 806, 871, 897, 898, 974, 975, 1056, 1060, 1136, 1137, 1138, 1139, 1167, 1228, 1229, 1238, 1241, 1245
- <schema name characteristic> • **975**, 1010
- <schema name clause> • 167, **567**, 568, 569, 1167, 1229
- <schema name list> • 193, **515**, 568, 736, 806, 978, 1237, 1238
- <schema path specification> • 81, **567**, 568, 569, 1192, 1237
- <schema procedure> • **715**
- <schema qualified name> • 82, 129, **161**, 162, 163, 166, 167, 529, 975, 1229
- <schema qualified routine name> • 88, **162**, 169, 541, 542, 543, 693, 715, 718, 721, 723, 724, 732, 733, 737, 1215
- <schema routine> • 106, 568, **715**, 718, 737, 1132, 1215
- schema-level routine • 88, 718, 732, 733, 735, 736, 740, 743, 833
- <schema-resolved user-defined type name> • **163**, 165, 179, 541, 542, 543, 674, 676, 691, 692, 693, 695, 697, 703, 705, 708, 712, 716, 719, 745, 746, 749, 753, 755, 758, 759, 761, 763, 779, 780, 781, 1187, 1189
- SCHEMA\_NAME • 146, 1126, 1136, 1137, 1138
- SCOPE • 148, 175, 465, 578, 610, 630
- scope • **7**, 46, 47, 55, 87, 99, 100, 136, 168, 169, 170, 179, 198, 199, 200, 202, 216, 218, 219, 241, 243, 244, 246, 329, 330, 358, 359, 367, 368, 372, 377, 387, 388, 465, 579, 580, 603, 609, 610, 611, 622, 623, 624, 630, 632, 635, 636, 668, 692, 723, 758, 768, 771, 773, 792, 793, 794, 795, 796, 798, 849, 855, 856, 857, 861, 866, 867, 868, 879, 886, 888, 986, 1035, 1037, 1039, 1041, 1043, 1044, 1046, 1050, 1053, 1059, 1076, 1271
- <scope clause> • **175**, 179, 183, 184, 216, 329, 330, 574, 579, 582, 609, 625, 629, 1191
- <scope option> • 163, **164**, 168, 170, 853, 1037, 1038, 1050
- SCOPE\_CATALOG • 146, 984, 986, 997, 1000, 1002, 1021
- SCOPE\_NAME • 146, 984, 986, 997, 1000, 1002, 1021
- SCOPE\_SCHEMA • 146, 984, 986, 997, 1000, 1002, 1021
- scoped table • 46, 246, 632, 771, 772, 773, 792, 793, 795, 796, 799
- SCROLL • 94, 100, 102, 148, 851, 854, 857, 904, 905, 1038
- scrollable • 94, 101, 102, 537, 852, 860, 1173, 1174, 1277, 1279
- SEARCH • 148, 393
- <search clause> • **393**, 394, 395, 396
- <search condition> • 33, 52, 64, 67, 69, 70, 71, 73, 75, 77, 78, 90, 125, 203, 205, 211, 214, 230, 270, 283, 329, 332, 335, 336, 339, 344, 352, 354, 355, **451**, 547, 551, 586, 588, 594, 595, 596, 599, 601, 602, 605, 610, 611, 612, 616, 619, 622, 635, 638, 639, 644, 646, 650, 654, 659, 661, 662, 663, 664, 666, 695, 712, 713, 740, 743, 747, 748, 753, 754, 771, 789, 792, 793, 794, 795, 796, 798, 799, 836, 868, 869, 870, 878, 879, 881, 882, 883, 885, 888, 889, 890, 891, 941, 1134, 1135, 1170, 1178, 1183, 1214, 1247, 1262
- search condition too long for information schema* • 602, 662, 1151
- <search or cycle clause> • 376, 380, **393**, 397, 1255
- <searched case> • **211**, 213, 1271
- <searched when clause> • **211**, 213, 214



- SECOND • 35, 38, 148, 158, 181, 259, 260, 464, 509, 510, 511, 590, 1002, 1006, 1011, 1232
- secondary effects • 59, 60
- <seconds fraction> • **155**, 157, 158, 159, 160, 511, 1176
- <seconds integer value> • **155**, 158, 511
- <seconds value> • **154**, 158
- SECTION • 146, 1056
- SECURITY • 146, 180, 203, 217, 218, 244, 280, 332, 540, 544, 630, 717, 725, 732, 733, 734, 735, 737, 867, 868, 875, 880, 887, 888, 1044, 1046, 1049, 1053
- SELECT • 59, 75, 76, 78, 120, 121, 148, 203, 209, 219, 243, 246, 247, 249, 308, 309, 311, 312, 327, 328, 329, 332, 336, 337, 341, 342, 343, 347, 351, 352, 366, 370, 380, 382, 386, 389, 394, 395, 396, 555, 556, 580, 581, 582, 588, 594, 595, 596, 599, 605, 616, 617, 626, 631, 632, 633, 669, 702, 771, 772, 773, 774, 776, 777, 779, 780, 781, 786, 789, 791, 792, 793, 794, 795, 796, 797, 798, 799, 862, 869, 880, 881, 883, 888, 889, 898, 920, 924, 931, 1134, 1135, 1214
- <select list> • 23, 29, 59, 63, 75, 78, 79, 198, 203, 205, 207, 208, 209, 230, 329, 336, 337, 342, 344, 346, 350, 359, **366**, 367, 369, 370, 371, 373, 374, 375, 379, 382, 385, 394, 432, 619, 629, 663, 794, 795, 796, 797, 799, 862, 863, 864, 919, 921, 928, 1017, 1018, 1028, 1214, 1218, 1261, 1264, 1265, 1270
- <select statement: single row> • 51, 62, 63, 107, 110, 113, 116, 117, 207, 209, 230, 329, 359, 794, 795, 796, 797, 799, 832, 833, 836, **862**, 863, 1133, 1249, 1250, 1270
- <select sublist> • 277, **366**, 368, 629, 919
- <select target list> • 209, **862**, 863, 864, 865, 1249
- SELF • 42, 146, 675, 677, 679, 680, 681, 682, 684, 685, 686, 688, 697, 698, 699, 700, 705, 1188
- <self-referencing column name> • 573, **574**, 578, 630
- <self-referencing column specification> • **573**, 577, 625, 627, 628
- <semicolon> • 140, **141**, 666, 811, 1055, 1057, 1078, 1100, 1105, 1111, 1119
- SENSITIVE • 94, 100, 102, 148, 851, 852, 1213
- sensitive • 102, 103, 150, 869, 871, 876, 881, 889, 903, 908, 910, 1227, 1229, 1239, 1240
- SEPARATE • 1088
- <separator> • **145**, 148, 149, 152, 153, 155, 156, 157, 254, 893, 1057, 1068, 1095, 1282
- SEQUENCE • 146, 571, 766, 768, 769, 779
- sequence • **7**
- <sequence column> • **393**, 394
- sequence generator • 53, 60, 81, 82, 83, 120, 121, 230, 231, 504, 505, 506, 507, 571, 575, 587, 615, 766, 767, 768, 769, 788, 795, 798
- <sequence generator cycle option> • 505, 506, 507, 766, **767**
- <sequence generator data type option> • **766**, 767, 1238
- <sequence generator definition> • 106, 568, **766**, 767, 831, 1133, 1212, 1238
- <sequence generator increment> • 505, 507, **766**
- <sequence generator increment by option> • 505, 507, **766**, 767
- sequence generator limit exceeded* • 504, 1146
- <sequence generator max value> • 505, 507, **766**, 1235
- <sequence generator maxvalue option> • 505, 507, **766**, 767, 1235
- <sequence generator min value> • 505, 507, 766, **767**, 1235
- <sequence generator minvalue option> • 505, 507, **766**, 767, 1235
- <sequence generator name> • **163**, 169, 171, 230, 231, 766, 767, 768, 769, 779, 780, 1211
- <sequence generator option> • **766**
- <sequence generator options> • **766**
- <sequence generator restart value> • 508, **768**
- <sequence generator start value> • 505, **766**
- <sequence generator start with option> • 505, **766**, 767
- SERIALIZABLE • 66, 124, 125, 126, 127, 146, 948, 949, 950, 1121, 1167
- serialization failure* • 126, 1150
- SERVER\_NAME • 146, 1126, 1139
- SESSION • 146, 969, 971
- <session characteristic> • **969**
- <session characteristic list> • **969**
- <session transaction characteristics> • **969**
- SESSION\_USER • 118, 148, 190, 191, 193, 194, 251, 589, 590, 591, 592, 902, 1123, 1170, 1231
- SET • 16, 69, 70, 127, 133, 148, 173, 176, 177, 311, 393, 448, 539, 567, 571, 597, 599, 604, 607, 612, 615, 642, 648, 650, 651, 779, 822, 835, 838, 839, 841, 842, 843, 878, 886, 888, 934, 935, 936, 937, 945, 947, 948, 951, 955, 964, 969, 971, 972, 973, 974, 975, 977, 978, 979, 980, 999, 1011, 1046, 1052, 1064, 1071, 1072, 1073, 1078, 1079, 1081, 1084, 1087, 1088, 1089, 1094, 1095, 1100, 1101, 1105, 1106, 1109, 1111, 1112, 1121, 1213, 1230
- <set catalog statement> • 86, 108, 129, 833, **974**, 1133, 1178, 1182
- <set clause> • 882, 890, **892**, 893, 894, 895, 896, 911, 1010, 1183, 1189
- <set clause list> • 329, 878, 879, 880, 882, 886, 887, 888, 890, **892**, 893, 894, 1046, 1047, 1052, 1053, 1222
- <set column default clause> • **606**, **607**, 1170
- <set connection statement> • 108, 128, 137, 832, 835, **964**, 965, 1121, 1133, 1182
- <set constraints mode statement> • 67, 68, 108, 832, **951**, 952, 1133, 1181
- <set descriptor information> • **999**

- <set descriptor statement> • 86, 109, 111, 833, 998, **999**, 1000, 1002, 1133, 1160, 1250
- <set domain default clause> • 641, **642**, 1180
- <set function specification> • 78, 79, 188, 189, 203, **205**, 206, 336, 344, 353, 355, 358, 365, 369, 370, 371, 372, 374, 379, 385, 386, 549, 550, 551, 601, 619, 663, 893, 1174, 1214, 1264, 1270
- <set function type> • **547**, 549
- <set header information> • **999**, 1002
- <set item information> • **999**, 1000
- <set local time zone statement> • 108, 129, 833, **973**, 1010, 1133, 1173
- <set names statement> • 86, 109, 130, 833, **977**, 1133, 1175, 1182
- set of involved grantees • 777, 784
- set of involved privilege descriptors • 777, 784
- set of overriding methods • 523, 741
- set of subject types • 165
- set of transitions • 135, 836, 913, 917, 920, 923, 926, 930, 940
- set operator • 380, 382, 384, 385, 389, 393
- <set path statement> • 86, 109, 130, 833, **978**, 1133, 1192, 1229
- <set predicate> • 401, 402, **448**, 483, 1201
- <set predicate part 2> • 212, **448**
- <set quantifier> • 23, 29, 209, 230, 345, 346, 353, 366, 367, 369, 370, 374, 375, 386, **547**, 548, 549, 557, 862, 863, 1174, 1184, 1208, 1218, 1248, 1260, 1263, 1264
- <set role statement> • 108, 118, 832, **972**, 1133, 1216
- <set schema statement> • 86, 108, 129, 130, 833, **975**, 976, 1133, 1182
- <set session characteristics statement> • 108, 124, 131, 833, **969**, 970, 1133, 1181, 1182
- <set session collation statement> • 109, 130, 833, **980**, 981, 1133, 1180, 1241
- <set session user identifier statement> • 108, 118, 832, **971**, 1010, 1133, 1170, 1241
- <set target> • **892**, 893, 1010
- <set target list> • **892**, 893, 1010
- <set time zone value> • **973**
- <set transaction statement> • 98, 108, 124, 832, 945, 946, **947**, 948, 1121, 1133, 1213, 1267
- <set transform group statement> • 109, 130, 833, **979**, 1004, 1133, 1199
- SETS • 63, 78, 146, 345, 347, 348, 349, 350, 353, 716, 720, 1157, 1217
- SIGN • 1088
- <sign> • **153**, 154, 159, 222, 223, 254, 293, 294, 1258
- signature • **9**
- <signed integer> • 31, **153**
- <signed numeric literal> • 152, **153**, 159, 220, 505, 507, 590, 766, 767, 768, 1235
- significant • 28, 30, 31, 35, 38, 39, 66, 103, 157, 220, 228, 255, 289, 295, 296, 297, 409, 464, 510, 511, 825, 869, 871, 876, 881, 889, 903, 908, 911, 1095, 1226, 1227, 1239, 1240
- SIMILAR • 20, 23, 148, 273, 283, 420, 421, 422, 423, 425, 1005, 1210
- <similar pattern> • **420**, 421, 422, 423, 479, 1009
- <similar predicate> • 23, 401, 402, **420**, 421, 425, 478, 479, 1009, 1207, 1210, 1234
- <similar predicate part 2> • 211, **420**
- SIMPLE • 68, 69, 146, 436, 597, 599, 933, 934, 935, 936
- simple • 36, 63, 207, 208, 214, 346, 347, 349, 359, 385, 1109, 1168, 1276
- <simple case> • **211**, 212, 213, 214, 1008, 1168, 1271
- <simple comment> • **145**, 149, 1011, 1267
- <simple comment introducer> • **145**, 149
- <simple Latin letter> • **139**, 161, 424
- <simple Latin lower case letter> • **139**, 142, 150, 424
- <simple Latin upper case letter> • **139**, 142, 150, 424, 1067, 1143, 1243
- <simple table> • 79, **376**, 379, 380, 381, 384, 385, 389, 391, 565, 1178, 1248
- simple table query • 385
- <simple target specification> • 136, **190**, 192, 193, 194, 472, 473, 985, 986, 996, 997, 1125, 1126, 1129, 1136, 1140, 1141, 1167, 1231
- <simple target specification 1> • **996**, 997, 998, 1250
- <simple target specification 2> • **996**, 997, 998, 1250
- <simple value specification> • 127, 163, 164, 168, 170, **190**, 191, 192, 193, 277, 472, 473, 535, 857, 858, 859, 864, 892, 894, 895, 896, 905, 949, 962, 985, 986, 993, 996, 999, 1003, 1004, 1012, 1036, 1126, 1194
- <simple value specification 1> • **999**, 1002
- <simple value specification 2> • **999**, 1000
- <simple when clause> • **211**, 213, 214, 1168
- simply dependent • 561
- simply underlying table • 57, 372, 373, 384, 866, 886, 909, 911, 1044, 1046, 1048, 1052
- simply updatable • 373, 380, 381, 626
- simply updatable derived table • 331
- single character specifier • 416, 417
- <single datetime field> • 179, **509**, 510, 511
- <single group specification> • **717**, 722, 806, 1058
- site • **7**
- SIZE • 146, 945, 949, 950
- sizing item • **7**
- SMALLINT • 14, 15, 30, 148, 174, 178, 182, 469, 474, 815, 820, 822, 840, 841, 843, 844, 845, 847, 984, 985, 1076, 1084, 1092, 1116, 1230

- <solidus> • **140**, 149, 254, 255, 293, 1006
- SOME • 65, 148, 430, 547, 549, 552, 557, 1204
- <some> • **430**, 431
- sort direction • 559
- <sort key> • 62, 208, 358, 359, 361, 362, 363, 385, 485, 550, 555, 556, **559**, 1011, 1264, 1265
- <sort specification> • 62, 359, 385, 550, 551, 555, 558, **559**, 560, 1175, 1236, 1249
- <sort specification list> • 62, 356, 360, 377, 385, 485, 548, 550, 555, 556, 557, 558, **559**, 1174, 1237, 1249, 1255
- sort table • 386, 390
- SOURCE • 146, 675
- source • 7, 20, 22, 37, 45, 50, 51, 60, 96, 148, 149, 156, 158, 159, 203, 206, 216, 217, 353, 374, 477, 548, 549, 586, 656, 713, 714, 744, 746, 747, 814, 877, 884, 919, 920, 1229, 1232
- <source character set specification> • **656**, 789
- <source data type> • **745**, 747
- source type • 40
- SPACE • 19, 25, 146, 421, 424, 540, 652
- <space> • 10, 18, 19, 87, 127, 139, **140**, 141, 148, 153, 154, 155, 160, 193, 220, 221, 222, 223, 270, 271, 284, 285, 408, 424, 455, 459, 460, 591, 825, 827, 828, 1058, 1095, 1138, 1204
- SPECIFIC • 148, 541, 571, 611, 620, 624, 637, 664, 676, 713, 716, 738, 744, 748, 751, 753, 762, 764, 765, 802, 1190
- <specific method name> • 40, 42, **676**, 679, 685, 686, 697, 698, 701, 702, 703, 706, 707, 708, 716, 718, 719, 720
- <specific method specification designator> • **708**
- <specific name> • 22, 45, 51, 89, 92, 93, 132, **162**, 169, 500, 502, 524, 529, 533, 534, 541, 571, 610, 620, 624, 637, 657, 664, 713, 716, 719, 720, 721, 733, 734, 735, 740, 743, 746, 747, 748, 749, 750, 753, 757, 759, 760, 762, 763, 764, 765, 802, 1022, 1138, 1139
- specific name text item • 528, 529
- <specific routine designator> • 41, 100, 526, **541**, 542, 543, 656, 657, 740, 743, 744, 745, 749, 750, 751, 755, 778, 779, 780, 781, 1036, 1037, 1038, 1188, 1189
- <specific type method> • 273, **274**, 275, 279, 281, 285, 286, 287, 1200
- SPECIFIC\_NAME • 146, 1126, 1138, 1139
- SPECIFICTYPE • 148, 274, 751
- specified type • 449, 450
- SQL • 9, 89, 90, 148, 180, 203, 217, 218, 244, 280, 332, 372, 513, 527, 531, 532, 534, 540, 544, 630, 677, 679, 680, 685, 687, 688, 698, 699, 701, 717, 720, 722, 725, 732, 733, 734, 735, 736, 741, 742, 751, 755, 761, 806, 867, 868, 875, 880, 887, 888, 993, 995, 996, 999, 1017, 1027, 1044, 1046, 1049, 1053, 1055, 1056, 1072, 1073, 1074, 1075, 1079, 1080, 1081, 1082, 1083, 1084, 1087, 1088, 1089, 1090, 1091, 1094, 1095, 1096, 1097, 1098, 1100, 1101, 1102, 1103, 1105, 1106, 1107, 1108, 1109, 1111, 1112, 1113, 1114, 1115, 1188, 1237, 1238
- <SQL argument> • **9**, 235, 379, 472, **516**, 517, 518, 519, 520, 535, 538, 1010, 1012, 1162, 1195
- <SQL argument list> • 90, 189, 235, 237, 239, 241, 244, 379, 472, **516**, 517, 518, 519, 520, 724, 1012
- <SQL condition> • **1067**, 1070, 1163, 1176
- <SQL connection statement> • 89, 137, 668, 725, 831, **832**, 833, 834, 835, 1119, 1120, 1121, 1139, 1235, 1237, 1238
- <SQL control statement> • 89, 112, 831, **832**, 833, 1003, 1004
- <SQL data change statement> • 668, 725, **832**, 1237, 1238
- <SQL data statement> • 725, 831, **832**, 834, 1238
- <SQL descriptor statement> • **833**
- <SQL diagnostics information> • **1125**, 1139
- <SQL diagnostics statement> • 98, 137, 831, **833**, 835, 836, 837
- <SQL dynamic data statement> • 123, **833**, 834, 1228
- <SQL dynamic statement> • 170, 668, 725, 831, **833**, 1023, 1027, 1237, 1238
- <SQL executable statement> • **831**
- SQL indicator argument • 528, 529, 531, 534
- <SQL language character> • 24, **139**, 177, 221, 222, 223, 568, 810, 1058, 1225, 1238
- <SQL language identifier> • 143, **161**, 163, 164, 168
- <SQL language identifier part> • **161**, 164
- <SQL language identifier start> • **161**
- SQL parameter • **9**
- <SQL parameter declaration> • 88, 89, 93, 650, 654, 679, 680, 682, 698, 700, 705, 706, **715**, 721, 722, 723, 726, 728, 729, 730, 731, 735, 737, 1188
- <SQL parameter declaration list> • 40, 41, 42, 198, 199, 367, 542, 676, 679, 680, 685, 686, 697, 698, 703, 704, **715**, 716, 719, 721, 722
- <SQL parameter name> • 92, **162**, 169, 198, 199, 367, 370, 371, 661, 679, 680, 683, 687, 688, 698, 704, 715, 719, 722, 723, 1063, 1139
- <SQL parameter reference> • 168, 190, 191, 192, 193, 200, 201, **204**, 370, 371, 520, 535, 574, 584, 601, 625, 668, 858, 859, 862, 864, 901, 902, 1023, 1027, 1215
- <SQL prefix> • **1055**, 1056, 1057, 1058
- <SQL procedure statement> • 66, 85, 89, 90, 95, 98, 113, 114, 137, 170, 292, 329, 610, 666, 717, 725, 733, 774, 811, 812, 813, **831**, 833, 834, 836, 934, 941, 1004, 1033, 1055, 1061, 1063, 1064, 1065, 1068, 1069, 1120, 1237, 1238, 1246, 1251
- SQL routine • **9**, 89, 90, 92, 132, 372, 523, 524, 526, 527, 537, 570, 571, 610, 612, 616, 620, 622, 635, 646, 650, 654, 659, 663, 664, 695, 712, 713, 720, 723, 725, 733, 740, 743, 747, 748, 753, 769, 797, 798, 799, 944, 1237, 1238

<SQL routine body> • 619, 663, **717**, 725, 734, 735, 738, 798, 799, 1222, 1223, 1235, 1238  
*SQL routine exception* • 525, 526, 527, 1138, 1150  
 <SQL routine spec> • 180, 203, 217, 218, 244, 280, 332, 540, 544, 630, **717**, 725, 732, 867, 868, 875, 880, 887, 888, 1044, 1046, 1049, 1053  
 <SQL schema definition statement> • **831**  
 <SQL schema manipulation statement> • **831**  
 <SQL schema statement> • 81, 82, 99, 136, 165, 168, 180, 203, 217, 218, 244, 280, 332, 517, 540, 544, 592, 605, 630, 668, 725, 732, **831**, 833, 835, 836, 849, 867, 868, 875, 880, 887, 888, 1003, 1035, 1044, 1046, 1049, 1050, 1053, 1119, 1121, 1237, 1238  
 <SQL session statement> • 668, 831, **832**, 1004, 1119, 1237  
 <SQL special character> • **139**, 144  
 <SQL statement name> • **11**, **163**, 171, 531, 1003, 1012, 1015, 1017, 1018, 1032, 1160, 1271  
 <SQL statement variable> • **11**, 111, 853, **1003**, 1004, 1034  
 <SQL terminal character> • **139**, 1229  
 <SQL terminator> • **1055**, 1056, 1057  
 <SQL transaction statement> • 89, 668, 725, 831, **832**, 1003, 1119, 1235, 1237, 1238  
 SQL-client module • 55, 56, 83, 84, 87, 91, 93, 99, 100, 104, 119, 123, 128, 129, 151, 156, 157, 193, 487, 525, 568, 639, 797, 805, 807, 810, 813, 834, 835, 849, 898, 955, 961, 1013, 1035, 1037, 1038, 1050, 1059, 1065, 1171, 1227, 1228, 1229, 1231, 1235, 1240, 1245  
 SQL-client module collations • 84  
 <SQL-client module definition> • 55, 83, 84, 85, 87, 91, 99, 100, 113, 114, 119, 129, 157, 164, 165, 166, 167, 170, 191, 192, 202, 518, 519, 525, 568, 580, 629, 639, 648, 649, 652, 657, 661, 667, 668, 684, 692, 718, 732, 733, 736, 767, **805**, 806, 807, 808, 809, 810, 811, 812, 813, 814, 849, 853, 897, 1013, 1015, 1017, 1032, 1033, 1035, 1037, 1038, 1059, 1065, 1163, 1164, 1227, 1235, 1251, 1271  
 <SQL-client module name> • 137, **162**, 169, 810, 814, 1059, 1246, 1251  
*SQL-client unable to establish SQL-connection* • 962, 1144  
 SQL-connection • xx, **7**, 108, 112, 113, 115, 116, 127, 128, 129, 169, 526, 530, 869, 871, 876, 881, 889, 908, 910, 947, 955, 961, 962, 964, 965, 966, 967, 1240, 1250  
 <SQL-data access indication> • 676, 680, 688, 698, **716**, **717**, 720, 732, 736, 740, 741, 742, 1188  
 SQL-environment • **7**  
 SQL-implementation • **7**  
 SQL-invoked constructor method • 43, 88, 517, 543, 679, 697, 699, 718, 724  
 SQL-invoked constructor methods • 43, 88, 89  
 <SQL-invoked function> • **7**, 8, 9, 10, 17, 44, 45, 50, 51, 88, 89, 90, 91, 92, 479, 485, 502, 503, 527, 528, 531,

535, 538, 542, 543, 682, 684, 687, 700, 705, 710, **715**, 718, 721, 723, 726, 731, 733, 734, 735, 755, 756, 759, 902, 944, 1162, 1274  
 SQL-invoked method • 40, 47, 88, 89, 90, 93, 169, 235, 241, 406, 516, 517, 522, 523, 531, 533, 534, 542, 543, 684, 691, 700, 705, 710, 711, 718, 724, 726, 727, 728, 729, 731, 734, 735, 741, 756, 759, 760, 1026  
 <SQL-invoked procedure> • **9**, 88, 90, 91, 92, 93, 101, 103, 132, 526, 527, 533, 538, 542, 543, **715**, 718, 720, 723, 728, 733, 741, 902, 903, 943, 1036, 1037, 1148, 1235, 1273  
 SQL-invoked regular function • 89, 90, 517, 542, 543, 720, 721, 724, 750  
 <SQL-invoked routine> • **8**, 9, 11, 22, 66, 82, 87, 88, 89, 90, 91, 92, 93, 94, 118, 120, 121, 130, 132, 169, 212, 244, 372, 385, 466, 515, 516, 517, 518, 519, 521, 525, 533, 541, 542, 543, 571, 602, 610, 620, 623, 624, 636, 657, 661, 664, 668, 687, 693, 712, 713, 714, **715**, 718, 720, 721, 722, 723, 724, 725, 731, 732, 733, 734, 735, 736, 737, 738, 740, 742, 743, 746, 748, 750, 751, 759, 760, 774, 780, 788, 794, 802, 831, 833, 868, 880, 888, 907, 941, 943, 1012, 1015, 1017, 1032, 1036, 1037, 1065, 1138, 1139, 1219, 1220, 1223, 1237, 1238, 1273, 1287  
 <SQL-path characteristic> • **978**, 1010  
 <SQL-server name> • 127, **163**, 168, 961, 962, 1139, 1228, 1240  
*SQL-server rejected establishment of SQL-connection* • 962, 1144  
 SQL-session • xx, **7**, 19, 36, 37, 55, 56, 87, 90, 91, 95, 96, 97, 100, 101, 108, 112, 113, 115, 116, 118, 119, 122, 123, 124, 125, 128, 129, 130, 131, 132, 137, 158, 160, 164, 165, 166, 167, 170, 192, 193, 225, 289, 292, 487, 518, 519, 524, 526, 537, 538, 568, 733, 735, 736, 812, 813, 834, 835, 836, 856, 858, 861, 867, 887, 897, 898, 903, 941, 945, 946, 948, 951, 955, 956, 957, 958, 961, 962, 963, 964, 965, 966, 969, 971, 972, 973, 974, 975, 977, 978, 979, 980, 981, 1001, 1025, 1029, 1037, 1038, 1050, 1120, 1121, 1155, 1227, 1228, 1229, 1235, 1236, 1241, 1245, 1246  
 SQL-session role name • 118, 129, 971, 972  
 SQL-session user identifier • 118, 119, 129, 538, 971, 1228, 1241  
 SQL-statement • **7**  
 SQL-transaction • xx, 9, 67, 68, 97, 102, 103, 104, 107, 110, 111, 112, 113, 115, 116, 122, 123, 124, 125, 126, 127, 128, 169, 398, 527, 530, 531, 835, 836, 869, 871, 876, 881, 889, 903, 908, 909, 910, 911, 945, 946, 947, 948, 951, 953, 955, 956, 957, 958, 961, 964, 972, 1121, 1122, 1135, 1136, 1149, 1227, 1228, 1239, 1240, 1246  
 SQL-update operation • 135  
 SQLEXCEPTION • 148, 1067, 1068, 1070  
 SQLSTATE • 90, 95, 96, 98, 148, 531, 728, 811, 812, 822, 823, 824, 838, 839, 841, 842, 844, 845, 846, 1059, 1060,

- 1062, 1065, 1067, 1068, 1069, 1070, 1076, 1129, 1136, 1143, 1144, 1152, 1163, 1243, 1246
- <SQLSTATE char> • **1067**, 1068
- <SQLSTATE class value> • **1067**, 1068, 1069, 1070
- SQLSTATE host parameter • **812**
- <SQLSTATE subclass value> • **1067**, 1068, 1069
- SQLWARNING • **148**, 1067, 1068, 1070
- SQRT • **148**, 258, 549, 554
- <square root> • **32**, 256, **258**, 260, 266, 1222
- ST-EC-ordered • **17**
- ST-FC-ordered • **17**
- ST-NC-ordered • **17**
- ST-ordered • **17**, 479, 482, 484, 486, 1187, 1188
- <standard character set name> • **539**, 540
- standard-defined classes • **1143**
- standing cursor • **99**, 100, 101, 102, 849, 856, 857, 861, 866, 886, 901, 902
- START • **148**, 257, 259, 266, 576, 766, 945, 946, 1213
- <start field> • **179**, 296, 464, **509**, 510, 511
- <start position> • **256**, 259, 265, 273, 274, **275**, 276, 277, 278, 280, 281, 284, 286, 1011, 1247
- <start transaction statement> • **107**, 110, 111, 112, 124, 832, **945**, 946, 1121, 1133, 1213
- start value • **53**, 60, 82, 575, 587
- STATE • **17**, 41, 45, **146**, 404, 407, 442, 685, 749, 751, 754
- <state category> • **749**, 750, 751
- state changes • **135**, 136, 834, 913, 919, 925, 934
- STATEMENT • **146**, 666, 668, 669, **1126**, 1140
- statement area • **97**
- statement completion unknown* • **128**, 1150
- <statement cursor> • **99**, 101, 853, **1036**
- statement execution context • **113**, **131**, **132**, **834**, **837**
- <statement information> • **1125**, 1140
- <statement information item> • **1125**, 1126, 1129, 1140
- <statement information item name> • **1125**, 1126, 1127, 1139, 1140, 1141, 1167, 1219
- statement information items • **97**, 1251
- <statement name> • **100**, **163**, 170, 853, 901, 1015, 1017, 1032, 1034, 1035, 1039, 1251
- <statement or declaration> • **1055**, 1065
- statement too long for information schema* • **669**, 1151
- statement-level trigger • **133**, **134**, **135**, 669
- STATIC • **40**, 42, 87, 93, **148**, 518, 531, **541**, 542, 543, 676, 679, 681, 682, 683, 685, 688, 689, 697, 700, 701, 703, 704, **708**, 709, 710, 711, **716**, 718, 719, 735, 751, 805, 807, 808, 1013, **1056**, 1163, 1188
- static method • **237**, 679, 681, 682, 697, 699, 700
- <static method invocation> • **89**, **188**, 189, **237**, 238, 695, 713, 740, 743, 792, 793, 794, 796, 797, 1187
- <static method selection> • **237**, 517, 518, 519, 521
- static SQL argument list • **218**, 219, 236, 237, 244, 517, 519, 534, 943
- static SQL-invoked method • **718**
- static SQL-invoked methods • **88**, 89
- <status parameter> • **95**, 96, 472, **811**, 812, 837, 1143
- STDDEV\_POP • **65**, **148**, 547, 548, 549, 557, 1222
- STDDEV\_SAMP • **65**, **148**, 547, 548, 549, 557, 1222
- stratum • **378**, 379, 383, 387
- string data, length mismatch* • **1146**
- string data, right truncation* • **221**, 222, 223, 224, 270, 271, 282, 284, 455, 456, 459, 460, 461, 1139, 1146, 1151
- <string length> • **273**, 274, **275**, 278, 280, 281, 286
- string types • **15**
- <string value expression> • **191**, 192, 193, 250, 252, 257, 259, 260, 261, 262, 266, **268**, 270, 280, 287, 421, 479, 1005, 1173, 1283
- <string value function> • **268**, **273**, 275, 280, 1005, 1283
- strings • **4**, **15**, 19, 23, 26, 28, 31, 98, 157, 408, 423, 424, 821, 822
- strong candidate key • **80**
- STRUCTURE • **147**, **1017**
- structured type • **7**, **8**, **9**, 13, 16, 17, 40, 41, 42, 43, 44, 46, 52, 53, 56, 58, 59, 121, 169, 179, 184, 198, 233, 239, 246, 368, 477, 577, 578, 580, 582, 609, 619, 627, 628, 632, 676, 677, 679, 684, 687, 688, 689, 690, 693, 695, 697, 712, 737, 738, 756, 759, 776, 777, 781, 786, 792, 813, 838, 877, 885, 893, 896, 1077, 1085, 1187, 1189, 1196
- structured types • **40**, 43, 52, 1155, 1156
- STYLE • **147**, 372, 527, 529, 531, 532, 534, 680, 698, 716, 722, 725, 730, 734, 736, 741, 742
- SUBCLASS\_ORIGIN • **147**, 1126, 1136, 1152, 1243
- subfield • **9**
- subject parameter • **88**, 89, 90, 523, 724
- subject routine • **89**, 90, 91, 94, 95, 212, 236, 237, 244, 252, 372, 385, 466, 472, 518, 519, 520, 521, 524, 526, 533, 534, 535, 538, 602, 661, 668, 695, 713, 725, 733, 736, 738, 740, 743, 792, 793, 796, 797, 798, 868, 880, 888, 943, 978, 1010, 1012, 1021, 1223, 1237, 1238
- subject table • **74**, 133, 134, 135, 136, 667, 669, 671, 794, 836, 866, 868, 874, 878, 886, 888, 913, 917, 920, 923, 926, 930, 938, 1044, 1046, **1048**, **1052**
- SUBMULTISET • **148**, **446**
- <submultiset predicate> • **251**, 372, 401, 402, **446**, 447, 483, 1202
- <submultiset predicate part 2> • **212**, **446**
- subordinate descriptor areas • **983**, 1030
- subordinate descriptors • **983**, 984, 1019
- <subquery> • **398**, 1267, 1270
- subrow • **59**, 869, 872, 881, 889, 913, 933

SUBSTRING • 20, 148, 261, 273, 275, 278, 279, 280, 1005  
*substring error* • 281, 286, 1146  
 substring specifier • 416, 418  
 SUBSTRING\_REGEX • 21, 148, 273, 1256  
 subtable • 58, 59, 68, 333, 578, 609, 610, 627, 628, 629, 869, 872, 881, 889, 925, 926  
 <subtable clause> • 58, 573, **574**, 577, 578, 580, 581, 582, 1193  
 subtable family • 59  
 subtype • **9**, 13, 41, 43, 44, 45, 46, 47, 48, 89, 234, 235, 450, 454, 458, 518, 519, 521, 682, 684, 693, 695, 697, 700, 701, 703, 705, 706, 725, 750, 814, 815, 821, 822, 823, 1225  
 <subtype clause> • 43, **674**, 677, 678, 684, 685, 687  
 subtype family • 43, 44, 45, 403, 449, 465  
 <subtype operand> • **233**  
 <subtype treatment> • 188, 189, **233**, 234, 1195  
 subview • 627, 629  
 <subview clause> • **625**, 627, 628, 629, 630, 633  
*successful completion* • 95, 96, 490, 537, 837, 1069, 1122, 1143, 1150  
 SUM • 64, 148, 547, 549, 552, 1231, 1236  
 superrow • 59, 68, 623, 913, 925, 933, 938  
 supertable • 58, 59, 68, 246, 332, 578, 581, 618, 623, 624, 628, 632, 633, 637, 771, 772, 773, 774, 792, 793, 795, 796, 799, 869, 880, 881, 888, 889, 913, 919, 925, 926  
 <supertable clause> • **574**  
 <supertable name> • **574**, 580  
 supertype • **8**, **9**, 13, 16, 41, 43, 44, 46, 47, 48, 216, 218, 233, 235, 237, 404, 465, 517, 678, 685, 687, 688, 692, 693, 695, 697, 700, 701, 703, 705, 712, 725, 749, 750, 797, 799, 800  
 <supertype name> • **674**, 678  
 superview • 627, 628  
 supporting privilege descriptor • 121  
 suppresses duplicates • 388  
 SYMMETRIC • 148, 411, 1218  
*syntax error or access rule violation* • 836, 905, 1011, 1013, 1033, 1034, 1037, 1039, 1120, 1137, 1143, 1150  
 SYSTEM • 148, 325, 574, 577, 582, 627, 675, 873, 877, 884, 1191  
 system-defined representation • 47, 581, 685  
 <system-generated representation> • **675**, 677  
 SYSTEM\_USER • 148, 190, 191, 193, 194, 251, 372, 589, 590, 591, 592, 902, 1123, 1170, 1231

## — T —

T • 143, 147, 151, 176, 1156, 1207, 1283  
 Feature T021, “BINARY and VARBINARY data types” • 160, 185, 1077, 1086, 1093, 1099, 1104, 1110, 1116, 1117, 1202, 1203, 1204

Feature T022, “Advanced support for BINARY and VARBINARY data types” • 287, 418, 1203, 1207  
 Feature T023, “Compound binary literals” • 160, 1203, 1204  
 Feature T024, “Spaces in binary literals” • 160, 1204  
 Feature T031, “BOOLEAN data type” • 160, 184, 252, 302, 317, 557, 1204  
 Feature T041, “Basic LOB data type support” • 160, 184, 738, 1076, 1085, 1092, 1099, 1104, 1110, 1116, 1202, 1204, 1205, 1206  
 Feature T042, “Extended LOB data type support” • 229, 287, 418, 419, 425, 479, 687, 1173, 1203, 1206, 1207  
 Feature T043, “Multiplier T” • 151, 1207  
 Feature T044, “Multiplier P” • 151, 1207  
 Feature T051, “Row types” • 171, 184, 187, 232, 317, 319, 374, 1207, 1208  
 Feature T052, “MAX and MIN for row types” • 558, 1208  
 Feature T053, “Explicit aliases for all-fields reference” • 374, 1208  
 Feature T061, “UCS support” • 185, 287, 435, 1154, 1208, 1209  
 Feature T071, “BIGINT data type” • 185, 1077, 1085, 1209  
 Feature T101, “Enhanced nullability determination” • 52, 300, 338, 371, 382, 383  
 Feature T111, “Updatable joins, unions, and columns” • 373, 385, 626, 633, 855, 870, 877, 885, 891, 1209  
 Feature T121, “WITH (excluding RECURSIVE) in query expression” • 170, 334, 391, 1209, 1210  
 Feature T122, “WITH (excluding RECURSIVE) in subquery” • 391, 1210  
 Feature T131, “Recursive query” • 391, 633, 1210  
 Feature T132, “Recursive query in subquery” • 391, 1210  
 Feature T141, “SIMILAR predicate” • 425, 1210  
 Feature T151, “DISTINCT predicate” • 443, 1210  
 Feature T152, “DISTINCT predicate with negation” • 443, 1210, 1211  
 Feature T171, “LIKE clause in table definition” • 582, 1211  
 Feature T172, “AS subquery clause in table definition” • 582, 1211  
 Feature T173, “Extended LIKE clause in table definition” • 583, 1211  
 Feature T174, “Identity columns” • 588, 615, 1211  
 Feature T175, “Generated columns” • 588, 1211  
 Feature T176, “Sequence generator support” • 171, 231, 767, 768, 769, 1211, 1212  
 Feature T191, “Referential action RESTRICT” • 600, 1212  
 Feature T201, “Comparable data types for referential constraints” • 600, 1212  
 Feature T211, “Basic trigger capability” • 334, 670, 673, 781, 1212  
 Feature T212, “Enhanced trigger capability” • 670, 1212

- Feature T213, "INSTEAD OF triggers" • 670, 1212, 1213
- Feature T231, "Sensitive cursors" • 852, 1183, 1213
- Feature T241, "START TRANSACTION statement" • 946, 1213
- Feature T251, "SET TRANSACTION statement: LOCAL option" • 948, 1213
- Feature T261, "Chained transactions" • 956, 959, 1213
- Feature T271, "Savepoints" • 170, 953, 954, 958, 1213, 1214
- Feature T272, "Enhanced savepoint management" • 527, 530, 734, 738, 1214
- Feature T281, "SELECT privilege with column granularity" • 781, 1214
- Feature T285, "Enhanced derived column names" • 375, 1214
- Feature T301, "Functional dependencies" • 354, 355, 365, 374, 998, 1214, 1215
- Feature T312, "OVERLAY function" • 287, 1215
- Feature T322, "Overloading of SQL-invoked functions and procedures" • 737, 1215
- Feature T323, "Explicit security for external routines" • 737, 1215
- Feature T324, "Explicit security for SQL routines" • 738, 1215
- Feature T325, "Qualified SQL parameter references" • 201, 374, 1215
- Feature T326, "Table functions" • 313, 334, 738, 1215, 1216
- Feature T331, "Basic roles" • 170, 782, 784, 785, 802, 972, 1216
- Feature T332, "Extended roles" • 194, 569, 592, 781, 782, 804, 1216, 1217
- Feature T351, "Bracketed comments" • 151, 1217
- Feature T431, "Extended grouping capabilities" • 206, 353, 1217
- Feature T432, "Nested and concatenated GROUPING SETS" • 353, 1217
- Feature T433, "Multiargument GROUPING function" • 206, 1217, 1218
- Feature T434, "GROUP BY DISTINCT" • 353, 1218
- Feature T441, "ABS and MOD functions" • 266, 1218
- Feature T461, "Symmetric BETWEEN predicate" • 411, 1218
- Feature T471, "Result sets return value" • 737, 852, 1038, 1218
- Feature T491, "LATERAL derived table" • 333, 1218
- Feature T501, "Enhanced EXISTS predicate" • 432, 1218
- Feature T511, "Transaction counts" • 1141, 1219
- Feature T551, "Optional key words for default syntax" • 391, 852, 1219
- Feature T561, "Holdable locators" • 899, 900, 1219
- Feature T571, "Array-returning external SQL-invoked functions" • 688, 737, 1219
- Feature T572, "Multiset-returning external SQL-invoked functions" • 689, 737, 1219, 1220
- Feature T581, "Regular expression substring function" • 287, 1220
- Feature T591, "UNIQUE constraints of possibly null columns" • 596, 1220
- Feature T601, "Local cursor references" • 171, 1220
- Feature T611, "Elementary OLAP operations" • 210, 365, 560, 1220
- Feature T612, "Advanced OLAP operations" • 171, 210, 267, 365, 558, 1220, 1221
- Feature T613, "Sampling" • 334, 1221
- Feature T621, "Enhanced numeric functions" • 266, 267, 557, 558, 1221, 1222
- Feature T641, "Multiple column assignment" • 896, 1222
- Feature T651, "SQL-schema statements in SQL routines" • 527, 738, 1222
- Feature T652, "SQL-dynamic statements in SQL routines" • 527, 738, 1222
- Feature T653, "SQL-schema statements in external routines" • 530, 738, 1222
- Feature T654, "SQL-dynamic statements in external routines" • 530, 738, 1222, 1223
- Feature T655, "Cyclically dependent routines" • 738, 1223
- TABLE • 148, 312, 326, 376, 380, 388, 395, 570, 573, 603, 611, 617, 620, 622, 623, 624, 637, 647, 664, 666, 667, 668, 716, 718, 748, 754, 779, 801, 802, 871, 897, 898, 1165
- table • 7
- table associated with • 59, 199, 200, 333, 368, 669
- <table commit action> • 573, 897
- <table constraint> • 57, 58, 67, 68, 69, 71, 169, 581, 593, 594, 598, 612, 618, 619, 624, 637, 647, 663, 792, 801, 1136, 1138
- <table constraint definition> • 573, 576, 577, 578, 581, 586, 587, 593, 594, 596, 599, 618, 647, 1265
- <table contents source> • 573, 574
- <table definition> • 54, 55, 58, 105, 165, 167, 230, 567, 573, 575, 580, 582, 584, 586, 587, 593, 595, 596, 597, 601, 605, 831, 1133, 1272
- <table element> • 573, 580
- <table element list> • 573, 575, 577, 579, 897
- <table expression> • 54, 61, 62, 78, 202, 207, 209, 322, 354, 357, 366, 369, 371, 372, 373, 379, 385, 394, 794, 795, 796, 797, 799, 862, 863, 915, 921, 928, 1270
- <table factor> • v, 76, 77, 199, 325, 329, 331, 333, 335, 336, 378, 379
- <table function column list> • 716, 718
- <table function column list element> • 716
- <table function derived table> • 325, 326, 334, 1216

<table name> • 54, 55, 56, 58, 81, 82, **161**, 164, 165, 167, 168, 175, 179, 326, 330, 331, 332, 561, 563, 564, 570, 573, 574, 575, 576, 578, 580, 581, 586, 593, 595, 596, 597, 599, 603, 604, 609, 610, 611, 612, 616, 618, 619, 620, 622, 623, 624, 625, 626, 628, 629, 630, 631, 635, 636, 637, 645, 666, 667, 668, 669, 671, 701, 706, 748, 754, 772, 779, 780, 781, 786, 787, 790, 794, 795, 798, 801, 802, 855, 866, 868, 871, 873, 877, 878, 879, 881, 886, 888, 889, 897, 916, 928, 931, 955, 1033, 1044, 1046, 1048, 1052, 1058, 1134, 1137, 1183, 1193, 1260, 1268

<table or query name> • 57, 58, 75, 76, 200, 202, 325, **326**, 329, 330, 331, 332, 368, 372, 376, 377, 380, 381, 384, 628, 855, 866, 868, 869, 881, 883, 884, 886, 888, 889, 1052, 1183, 1260

<table primary> • v, 66, 76, **325**, 326, 329, 331, 332, 333, 336, 379, 380

table privilege descriptor • 120, 564, 565, 566

<table reference> • v, 59, 66, 74, 77, 202, 323, 324, **325**, 329, 330, 331, 332, 333, 334, 335, 336, 372, 373, 378, 379, 380, 387, 394, 563, 601, 626, 628, 631, 725, 772, 788, 789, 791, 792, 793, 794, 796, 797, 798, 799, 866, 869, 878, 879, 881, 883, 884, 886, 889, 915, 916, 921, 928, 1044, 1046, 1048, 1052, 1134, 1183, 1194, 1227, 1260, 1268, 1270

<table reference list> • 77, **323**, 324, 378

<table row value expression> • **318**, 319, 320

<table scope> • **573**, 577, 579, 582, 1176

table specified by the • 330

<table subquery> • 230, 251, 301, 312, 313, 325, 326, 367, 376, 380, **398**, 412, 430, 432, 433, 436, 437, 438, 478, 479, 481, 485, 574, 576, 1009, 1218, 1219, 1258, 1260, 1261, 1262

<table value constructor> • v, 54, 74, 317, **320**, 321, 376, 381, 384, 391, 412, 565, 875, 919, 1008, 1178, 1247

<table value constructor by query> • **312**, 313, 1215

table/method privilege descriptor • 121

TABLE\_NAME • 147, 1126, 1136, 1137, 1138

TABLESAMPLE • 148, 325

target • 7, 31, 33, 37, 50, 51, 96, 97, 192, 193, 216, 217, 234, 277, 454, 458, 657, 713, 714, 744, 746, 747, 859, 865, 1027, 1028, 1041, 1065, 1147, 1149, 1232, 1234, 1249, 1250

<target array element specification> • 190, **191**, 192, 194, 520, 535, 858, 859, 862, 863, 864, 1194

<target array reference> • **191**, 192

<target character set specification> • **656**, 789

<target data type> • 50, **745**, 747

<target specification> • 136, **190**, 192, 194, 472, 473, 516, 520, 535, 836, 857, 858, 859, 862, 863, 864, 865, 1012, 1027, 1028, 1029, 1030, 1231, 1241, 1273

<target subtype> • **233**, 234, 1195

<target table> • 794, 798, **866**, 867, 868, 869, 870, 871, 872, 873, 878, 879, 880, 881, 882, 883, 885, 886, 888, 889, 890, 891, 892, 1044, 1046, 1048, 1052, 1134, 1194, 1209

*target table disagrees with cursor specification* • 1044, 1046, 1150

TEMPORARY • 54, 55, 147, 573, 579, 580, 581, 897

<temporary table declaration> • 55, 84, 106, 112, 114, 116, 129, 164, 202, 524, 525, 585, 587, 601, 805, **897**, 898, 1055, 1058, 1061, 1065, 1119, 1176, 1235

<term> • **254**, 293, 294

THEN • 148, 208, 211, 212, 213, 308, 309, 311, 327, 395, 878, 1140

TIES • 147, 357, 365

TIME • 14, 15, 34, 35, 37, 98, 148, 153, 157, 158, 175, 179, 180, 183, 224, 225, 226, 227, 259, 288, 289, 291, 292, 296, 338, 370, 371, 470, 474, 477, 840, 842, 843, 844, 846, 847, 973, 1001, 1002, 1009, 1011, 1233

<time fractional seconds precision> • 34, 35, **175**, 179, 181, 183, 260, 498, 613, 838, 1231

<time interval> • **154**

<time literal> • **153**, 157, 158, 159, 160, 1176, 1268

<time precision> • **175**, 179, 181, 184, 225, 291, 292, 984, 1021, 1177, 1231

<time string> • 144, **153**

time types • 15

<time value> • **154**, 159, 1268, 1269

<time zone> • 251, **288**, 289, 290, 1173

time zone displacement • 35, 36, 37, 39, 129, 131, 160, 225, 226, 261, 289, 290, 292, 464, 524, 973, 1146, 1228, 1235

<time zone field> • 34, **257**, 259, 266, 1173

<time zone interval> • **154**, 157, 158, 159, 160, 1172, 1268, 1269

<time zone specifier> • **288**, 289

TIMESTAMP • 14, 15, 34, 35, 37, 98, 148, 153, 158, 175, 179, 180, 183, 224, 225, 226, 227, 259, 288, 291, 292, 296, 338, 370, 371, 470, 475, 477, 840, 842, 843, 845, 846, 847, 1002, 1009, 1011

<timestamp literal> • **153**, 158, 159, 160, 1176, 1269

<timestamp precision> • **175**, 179, 181, 184, 226, 227, 291, 292, 984, 1021, 1177, 1231

<timestamp string> • 144, 148, **153**

timestamp types • 15

TIMEZONE\_HOUR • 148, 181, 257, 261

TIMEZONE\_MINUTE • 148, 181, 257

TO • 20, 35, 129, 148, 228, 283, 289, 295, 296, 393, 409, 420, 421, 422, 423, 425, 464, 509, 510, 539, 582, 633, 648, 656, 687, 755, 761, 762, 772, 773, 774, 776, 777, 783, 835, 957, 961, 973, 1002, 1006, 1010, 1011, 1067, 1068, 1069, 1070, 1121

<to sql> • **755**, 756, 757, 759, 760, 761



- <to sql function> • **755**, 756, 759
- to-sql function • 45, 93, 502, 503, 533, 534, 727, 728, 732, 733, 734, 755, 759, 760, 761, 762, 763, 764, 765, 1026
- to-sql function associated with i-th SQL parameter • 729
- to-sql function associated with the i-th SQL parameter • 730, 731
- <token> • **143**, 148, 149, 1058, 1282
- too many* • 953, 1150
- TOP\_LEVEL\_COUNT • 147, 996, 999, 1018
- TRAILING • 148, 274, 285, 287
- TRANSACTION • 127, 147, 945, 946, 947, 948, 969, 1213
- <transaction access mode> • 948, **949**, 969, 970, 1267
- <transaction characteristics> • 945, 947, 948, **949**, 1167
- <transaction mode> • **949**, 969, 1266
- transaction resolution unknown* • 128, 1144
- transaction rollback* • 126, 128, 955, 1136, 1137, 1150, 1152
- transaction-initiating • 110, 111, 112, 123, 127, 947, 964, 1120, 1121
- TRANSACTION\_ACTIVE • 147, 1125, 1135, 1136, 1141, 1219
- TRANSACTIONS\_COMMITTED • 147, 1125, 1135, 1141, 1219
- TRANSACTIONS\_ROLLED\_BACK • 147, 1125, 1135, 1141, 1219
- <transcoding> • **6**, 22, **273**, 275, 276, 277, 281, 285, 287, 1180, 1233
- <transcoding name> • **163**, 166, 169, 171, 273, 277, 285, 1180, 1229
- TRANSFORM • 147, 686, 713, 717, 744, 755, 758, 763, 979
- <transform definition> • 45, 106, 568, **755**, 757, 831, 1133, 1199
- transform descriptor • 41, 45, 500, 502, 685, 712, 741, 743, 744, 755, 759, 761, 763
- <transform element> • **755**
- <transform element list> • **755**, 756, 759
- transform functions • 45, 130, 755, 763, 979, 1004
- <transform group> • **755**
- <transform group characteristic> • **979**, 1010
- transform group descriptors • 45
- <transform group element> • **763**
- <transform group specification> • **717**, 722, 737, 805, 1056, 1199
- transform groups • 45, 712, 758
- <transform kind> • **761**
- TRANSFORMS • 147, 755, 758, 763
- Transforms • 45
- <transforms to be dropped> • **763**
- transition table • 74, 133, 134, 135, 136, 669, 866, 868, 874, 879, 886, 888, 940, 1048, 1052
- <transition table name> • 54, 56, 58, 326, 330, 332, 334, 564, **666**, 1212
- <transition table or variable> • **666**
- <transition table or variable list> • **666**, 667
- transition variable • 134, 136, 192, 193, 200, 520, 535, 668, 669, 926, 940, 941
- TRANSLATE • 148, 273
- TRANSLATE\_REGEX • 21, 148, 274, 1256
- TRANSLATION • 148, 571, 656, 659, 772, 779
- transliteration • **9**
- <transliteration definition> • 105, 166, 567, **656**, 657, 658, 772, 789, 831, 1133, 1180, 1237
- <transliteration name> • 22, **163**, 166, 169, 171, 274, 277, 280, 284, 571, 639, 656, 657, 659, 772, 779, 780, 789, 1180
- <transliteration routine> • **656**, 657
- <transliteration source> • **656**, 657
- TREAT • 148, 233
- TRIGGER • 120, 148, 571, 582, 617, 621, 624, 632, 637, 664, 666, 668, 671, 748, 754, 779, 780, 781, 794, 801, 1212
- <trigger action time> • 133, 134, 200, **666**, 669, 670, 1213
- <trigger column list> • 136, **666**, 667, 668, 669
- <trigger definition> • 56, 106, 134, 135, 136, 200, 567, **666**, 667, 668, 669, 670, 831, 1133, 1212, 1237
- trigger deletable • 54, 55, 57, 631, 667, 670, 672, 868, 917
- <trigger event> • 133, 134, 135, 136, **666**, 667, 668, 669, 670, 836, 913, 917, 920, 923, 926, 930, 940
- trigger execution contexts • 131, 132, 135
- trigger insertable-into • 55
- <trigger name> • **162**, 169, 571, 666, 667, 668, 669, 671, 748, 754, 801, 1137
- trigger updatable • 54
- TRIGGER\_CATALOG • 147, 1126, 1137
- TRIGGER\_NAME • 147, 1126, 1137
- TRIGGER\_SCHEMA • 147, 1126, 1137
- <triggered action> • 87, 90, 118, 133, 134, 135, 136, 292, 605, 610, 616, 617, 619, 622, 635, 654, 659, 663, **666**, 668, 669, 670, 740, 743, 753, 871, 881, 882, 890, 940, 1012, 1212, 1237
- triggered action column set • 669
- triggered action exception* • 941, 1137, 1150
- triggered data change violation* • 937, 1136, 1137, 1150
- <triggered SQL statement> • 134, 136, 520, **666**, 668, 669, 794, 795, 796, 868, 875, 888, 940, 1237
- <triggered when clause> • **666**, 668
- TRIM • 148, 224, 225, 226, 227, 228, 274, 275, 278, 280, 961, 962, 971, 972, 974, 975, 977, 978, 979, 980, 993, 1012, 1036
- <trim character> • **274**, 278, 285
- trim error* • 285, 286, 1146

<trim function> • 22, 29, 273, **274**, 275, 278, 281, 285, 1259, 1283  
 <trim octet> • **275**, 280, 286  
 <trim operands> • **274**  
 <trim source> • **274**, 278, 280, 285, 286  
 <trim specification> • **274**, 275, 278, 280, 285, 287  
 TRUE • 148, 155, 160, 222, 224, 299, 301, 395, 407, 751, 828, 868  
 TRUNCATE • 148, 871, 1256  
 <truncate table statement> • 60, 107, 111, 114, 115, 117, 832, **871**, 872, 1003, 1119, 1133, 1168, 1239  
 <truth value> • **299**, 300, 303, 1177  
 TYPE • 147, 571, 612, 674, 692, 712, 714, 717, 779, 802, 979, 984, 985, 986, 997, 998, 1000, 1001, 1002, 1019, 1020, 1021, 1024, 1025, 1030, 1031, 1072, 1073, 1074, 1075, 1079, 1080, 1081, 1082, 1083, 1084, 1087, 1088, 1089, 1090, 1091, 1094, 1095, 1096, 1097, 1098, 1100, 1101, 1102, 1103, 1105, 1106, 1107, 1108, 1109, 1111, 1112, 1113, 1114, 1115, 1241, 1250  
 type designators • 14, 468  
 <type list> • **449**  
 type of the method • 88, 89  
 type precedence list • 44, 51, 90, 216, 217, 468, 470, 471, 523, 747  
 type precedence lists • 14, 41, 90  
 <type predicate> • 401, 402, **449**, 450, 1195  
 <type predicate part 2> • 212, **449**  
 type-preserving function • 10, 42, 88, 92, 521, 533, 535, 723, 734, 735, 759  
 typed table • 46, 58, 330, 332, 579, 793, 796, 799, 866, 1048  
 <typed table clause> • **573**, 577, 580, 581  
 <typed table element> • **573**  
 <typed table element list> • **573**, 579

## — U —

U • 144, 152  
 UCS • 259, 276, 277, 493, 494, 496, 497, 1234  
 UCS string • 28  
 UDT-EC-ordered • 17, 481, 484, 485  
 UDT-FC-ordered • 17  
 UDT-NC-ordered • 17, 481, 484, 485, 750  
 UESCAPE • 20, 144, 148, 149  
 unary group aggregate function • 64  
 unassigned • 28, 591, 876, 877, 884, 919, 925, 926  
 unaugmented SQL parameter declaration list • 680, 681, 683, 697, 699, 704, 718  
 UNBOUNDED • 147, 208, 357, 358, 361, 362, 363, 364  
 UNCOMMITTED • 124, 125, 127, 147, 948, 949  
 undefined • 463, 998, 1000, 1004, 1005, 1007, 1011, 1250

*undefined DATA value* • 998, 1147  
 UNDER • 44, 59, 120, 121, 147, 574, 580, 625, 630, 633, 674, 684, 688, 692, 774, 779, 780, 781, 792, 797, 1187, 1193  
 under privilege descriptor • 121  
 underlying columns • 381, 384, 385, 627, 629, 631, 772, 927  
 underlying table • 57, 59, 373, 381, 564, 592, 627, 668, 774, 855, 866, 870, 873, 877, 879, 884, 885, 886, 890, 891, 896, 908, 909, 910, 911, 924, 931, 1044, 1046, 1048, 1052, 1183  
 <underscore> • 22, 23, 140, **141**, 152, 161, 164, 416, 417, 418, 421, 422, 424, 821, 1260  
 <Unicode 4 digit escape value> • **144**, 149  
 <Unicode 6 digit escape value> • **144**, 149  
 <Unicode character escape value> • **144**, 149  
 <Unicode character string literal> • 143, 148, **152**, 155, 156, 160, 1172  
 <Unicode delimited identifier> • 143, **144**, 148, 149, 150, 151, 161, 1171  
 <Unicode delimiter body> • **144**, 149, 150  
 <Unicode escape character> • **144**, 149, 1229  
 <Unicode escape specifier> • **144**, 149, 152  
 <Unicode escape value> • **144**, 149, 152, 156  
 <Unicode identifier part> • **144**, 149  
 <Unicode representation> • **152**, 155  
 UNION • 23, 29, 50, 72, 79, 80, 148, 247, 251, 252, 308, 309, 310, 328, 341, 342, 352, 376, 378, 380, 381, 382, 383, 384, 385, 388, 389, 390, 391, 408, 481, 483, 555, 633, 915, 927, 1201, 1219, 1248, 1267, 1272  
 UNIQUE • 58, 68, 148, 433, 436, 437, 438, 578, 595, 596, 1169, 1202, 1220  
 <unique column list> • 68, 481, 577, **595**, 596, 598, 1220, 1266  
 unique columns • 68, 69, 595, 596, 598, 619  
 <unique constraint definition> • 23, 29, 481, 577, 593, **595**, 596, 598, 616, 1266  
 unique matching row • 933  
 <unique predicate> • 401, 402, **433**, 481, 1169  
 <unique specification> • 584, 586, **595**, 596, 1265  
 Universal Character Sets • 18  
 UNKNOWN • 148, 155, 160, 299  
 UNNAMED • 147, 997, 999, 1019, 1020, 1250  
 UNNEST • 148, 219, 249, 308, 309, 311, 326  
 <unqualified schema name> • **161**, 164, 166, 167, 168, 193, 286, 568, 975, 1136, 1137, 1138, 1139  
 <unquoted date string> • 153, **154**, 224  
 <unquoted interval string> • **154**, 222, 223, 228  
 <unquoted time string> • 153, **154**, 225, 1268  
 <unquoted timestamp string> • 153, **154**, 226, 227, 1269

<unsigned integer> • **153**, 155, 159, 160, 174, 175, 176, 221, 223, 377, 387, 420, 509, 717, 1067, 1068, 1069, 1070, 1176, 1258  
 <unsigned literal> • **152**, 190  
 <unsigned numeric literal> • 143, 152, **153**  
 <unsigned value specification> • 188, 189, **190**, 192, 200, 357, 358, 361, 362, 363, 364  
 unsupported • 790  
*unterminated C string* • 825, 1146  
 <updatability clause> • 850, **853**, 854, 855, 887, 904, 1047, 1053, 1123, 1184  
 updatable • 52, 53, 54, 55, 57, 59, 60, 101, 102, 331, 373, 380, 381, 384, 385, 563, 564, 565, 566, 581, 587, 626, 629, 631, 632, 633, 667, 670, 672, 673, 774, 854, 855, 866, 868, 870, 874, 879, 880, 886, 888, 892, 893, 904, 930, 1040, 1044, 1046, 1048, 1052, 1209, 1267  
 updatable column • 54, 331, 385, 629  
 updatable columns • 384  
 updatable derived table • 331  
 UPDATE • 69, 70, 101, 120, 121, 134, 135, 136, 561, 562, 564, 565, 566, 581, 582, 597, 599, 604, 605, 617, 666, 667, 670, 672, 777, 779, 780, 781, 787, 795, 798, 850, 853, 854, 855, 878, 880, 886, 887, 888, 889, 898, 926, 930, 940, 1046, 1052, 1053, 1184, 1209, 1237  
 UPDATE • 148  
 update INSTEAD OF trigger • 133  
 update object column • 880  
 <update rule> • 69, 70, **597**, 599, 600, 933, 1180  
 <update source> • 230, 231, 794, 795, 796, 798, 799, 882, 890, **892**, 893, 894, 895, 896, 911, 1010, 1183, 1189  
 <update statement: positioned> • 51, 60, 94, 102, 103, 107, 111, 114, 116, 117, 231, 795, 798, 832, 870, 872, 882, **886**, 887, 890, 892, 909, 911, 938, 1133, 1184, 1264, 1265  
 <update statement: searched> • 51, 60, 107, 110, 114, 115, 117, 231, 604, 794, 795, 796, 798, 799, 832, **888**, 891, 892, 909, 911, 1003, 1119, 1133, 1134, 1135, 1183, 1209, 1251, 1264  
 <update target> • **892**, 894, 895, 896, 1010, 1189, 1194  
 update trigger • 133  
 update value • 895, 896  
 UPPER • 22, 148, 273, 283, 421, 424  
 <upper limit> • **420**, 423  
 USAGE • 120, 121, 147, 180, 217, 230, 231, 280, 539, 540, 544, 580, 587, 629, 639, 647, 648, 649, 651, 652, 653, 655, 657, 659, 688, 767, 769, 772, 779, 780, 781, 789, 792, 793, 794, 795, 796, 797, 798, 1087, 1088, 1090, 1091, 1189, 1238, 1242  
 usage privilege descriptor • 121  
 usage-dependent • 16, 792, 793, 794, 795, 796, 797, 798  
 USER • 148, 190, 191, 194, 574, 577, 589, 592, 628, 631, 873, 875, 877, 879, 884, 961, 1170

<user identifier> • 121, 161, **163**, 169, 961  
 <user-defined cast definition> • 51, 106, 568, **745**, 746, 831, 1133, 1195  
 <user-defined character set name> • **539**, 540  
 <user-defined ordering definition> • 106, 478, 479, 485, 568, **749**, 752, 831, 1134, 1200  
 <user-defined representation> • 42, 47, 581, 628, **675**, 677, 678, 685, 687  
 user-defined type • **10**, 826  
 <user-defined type body> • **674**, 676  
 <user-defined type definition> • 43, 44, 47, 53, 88, 106, 179, 567, **674**, 676, 684, 689, 691, 693, 694, 831, 1134, 1189, 1273  
 user-defined type locator parameter • 812, 826, 829, 899  
 user-defined type locator variable • 1075, 1083, 1091, 1097, 1103, 1108, 1114  
 <user-defined type name> • 40, 41, 84, 89, **163**, 165, 166, 169, 175, 179, 180, 233, 237, 286, 449, 519, 520, 523, 571, 577, 580, 627, 684, 723, 744, 802, 986, 1058, 1061, 1063  
 <user-defined type option> • **674**  
 <user-defined type option list> • **674**, 689, 1190  
 <user-defined type specification> • **449**  
 <user-defined type value expression> • **250**, 252, 274, 286  
 user-defined types • 13  
 USER\_DEFINED\_TYPE\_CATALOG • 147, 984, 986, 997, 1000, 1002, 1021  
 USER\_DEFINED\_TYPE\_CODE • 147, 997, 999, 1021  
 USER\_DEFINED\_TYPE\_NAME • 147, 984, 986, 997, 1000, 1002, 1021  
 USER\_DEFINED\_TYPE\_SCHEMA • 147, 984, 986, 997, 1000, 1002, 1021  
 USING • 148, 256, 257, 261, 265, 273, 274, 276, 281, 284, 335, 393, 478, 479, 675, 878, 1017, 1023  
 <using argument> • **1023**, 1025  
 <using arguments> • **1023**, 1025  
*using clause does not match dynamic parameter specifications* • 1023, 1024, 1147  
*using clause does not match target specifications* • 1027, 1028, 1147  
*using clause required for dynamic parameters* • 1033, 1039, 1147  
*using clause required for result fields* • 1033, 1147  
 <using descriptor> • **1017**, 1023  
 <using input descriptor> • **1023**, 1024

## — V —

vacant • 97  
 valid • 10, 11, 39, 51, 67, 97, 129, 131, 168, 215, 217, 228, 361, 810, 814, 821, 863, 899, 900, 956, 957, 958, 984, 985, 1024, 1028, 1059, 1065, 1068, 1072, 1075, 1076,

1081, 1084, 1089, 1091, 1095, 1098, 1101, 1103, 1109, 1112, 1115, 1229

VALUE • 52, 71, 148, 190, 192, 194, 230, 299, 300, 588, 595, 596, 873, 875, 877, 879, 884, 996, 999, 1091, 1168, 1202, 1242

<value expression> • 13, 59, 61, 64, 65, 66, 72, 79, 188, 189, 200, 205, 208, 211, 212, 213, 215, 216, 217, 218, 219, 233, **250**, 251, 252, 277, 292, 301, 306, 312, 315, 320, 336, 338, 344, 350, 352, 366, 367, 368, 369, 370, 371, 374, 375, 379, 385, 386, 393, 404, 405, 409, 481, 485, 516, 519, 520, 522, 547, 548, 549, 550, 551, 552, 555, 556, 557, 558, 559, 584, 687, 688, 747, 794, 795, 796, 798, 799, 833, 878, 892, 893, 896, 921, 928, 944, 1004, 1006, 1007, 1008, 1010, 1022, 1139, 1174, 1175, 1177, 1183, 1204, 1208, 1214, 1236, 1247, 1264, 1265, 1272

<value expression primary> • **188**, 189, 200, 232, 235, 241, 244, 250, 252, 254, 268, 288, 289, 293, 294, 304, 308, 309, 366, 367, 368, 610, 622, 632, 635, 771, 773, 792, 793, 794, 795, 796, 798, 1006, 1258, 1272, 1273

<value specification> • **190**, 191, 192, 251, 267, 277, 287, 413, 419, 427, 472, 473, 868, 875, 888, 902, 971, 972, 974, 975, 977, 978, 979, 980, 1010, 1012, 1119, 1123, 1169, 1177, 1185, 1231, 1242, 1273

VALUES • 59, 148, 320, 328, 394, 412, 555, 873, 874, 875, 877, 878, 1168

VAR\_POP • 65, 148, 547, 548, 549, 552, 557, 1222, 1236

VAR\_SAMP • 65, 148, 547, 548, 549, 552, 557, 1222, 1236

VARIABLE • 148, 174, 176, 474, 1072, 1073, 1079, 1082, 1084, 1094, 1096, 1101, 1102, 1111, 1114, 1256

VARCHAR • 148, 173, 176, 474, 1079, 1081, 1084, 1100

variable-length • **10**

VARYING • 14, 18, 28, 98, 148, 173, 174, 176, 177, 178, 182, 186, 228, 468, 469, 474, 585, 822, 825, 827, 828, 829, 838, 839, 840, 841, 842, 843, 844, 845, 846, 984, 985, 989, 1001, 1005, 1006, 1009, 1010, 1011, 1031, 1079, 1081, 1084, 1101, 1111, 1112, 1113, 1115, 1139

<vertical bar> • 23, 140, **141**, 420, 422, 423

VIEW • 147, 571, 611, 620, 624, 625, 626, 635, 636, 637, 664, 748, 754, 801

view • 54, 55, 56, 57, 58, 59, 81, 90, 119, 120, 121, 331, 561, 562, 563, 564, 565, 566, 570, 605, 610, 611, 612, 616, 619, 620, 622, 624, 627, 628, 630, 631, 632, 633, 635, 636, 646, 650, 654, 659, 663, 664, 668, 669, 671, 672, 673, 695, 712, 713, 740, 743, 747, 748, 753, 754, 772, 774, 775, 787, 792, 800, 801, 873, 879, 898, 918, 923, 924, 930, 931, 1137, 1237, 1270, 1276, 1278

<view column list> • **625**, 626, 630, 631, 788, 789, 791

<view column option> • **625**, 629

view component • 54, 565

view component column privilege descriptor • 120

view component table privilege descriptor • 120

<view definition> • 54, 58, 60, 105, 165, 561, 563, 567, **625**, 626, 627, 629, 630, 633, 634, 788, 789, 791, 831, 1134, 1181, 1186, 1210, 1269, 1272

<view element> • **625**

<view element list> • **625**, 629

view privilege dependency descriptor • 121

<view specification> • **625**

visible • 55, 103, 118, 134, 591, 869, 871, 876, 881, 889, 903, 908, 910, 1228, 1239, 1240, 1246

## — W —

warning • 95, 96, 123, 222, 223, 224, 284, 455, 456, 457, 490, 531, 536, 537, 552, 553, 555, 592, 602, 633, 662, 669, 777, 802, 837, 870, 872, 882, 890, 909, 911, 938, 948, 967, 1019, 1122, 1136, 1138, 1139, 1143, 1150

WHEN • 148, 208, 211, 212, 213, 308, 309, 311, 327, 395, 666, 878, 1140

<when operand> • **211**, 212, 213, 214, 1008, 1168

<when operand list> • **211**, 213, 214, 1168

WHENEVER • 148, 1067

WHERE • 78, 148, 243, 247, 328, 339, 344, 396, 547, 555, 556, 596, 866, 868, 883, 886, 888, 1044, 1046, 1048, 1052, 1134, 1135

<where clause> • v, 77, 202, 209, 322, 329, **344**, 346, 352, 359, 369, 373, 379, 1174, 1267, 1270

<white space> • **10**, **145**, 148, 149

WHITESPACE • 421, 424

<width bucket bound 1> • **258**, 264

<width bucket bound 2> • **258**, 264

<width bucket count> • **258**, 260, 264

<width bucket function> • 32, 256, **258**, 260, 264, 267, 1221

<width bucket operand> • **258**, 264

WIDTH\_BUCKET • 148, 258

WINDOW • 148, 209, 356

window • 61, 62, 63, 64, 66, 207, 208, 209, 210, 322, 356, 358, 359, 360, 361, 362, 363, 364, 365, 549, 1145, 1221, 1245, 1248

<window clause> • 63, 203, 205, 209, 322, 329, **356**, 357, 358, 359, 360, 365, 369, 370, 371, 1214, 1221, 1247

<window definition> • 61, **356**, 357, 359

<window definition list> • **356**

window frame • 61

<window frame between> • 356, **357**, 358

<window frame bound> • **357**

<window frame bound 1> • **357**, 358, 361

<window frame bound 2> • **357**, 358, 361

<window frame clause> • 62, **356**, 358, 360

<window frame exclusion> • 61, 356, **357**, 365, 1221

<window frame extent> • **356**, 358

<window frame following> • **357**, 358, 361, 363, 364, 1010

<window frame preceding> • **357**, 358, 361, 362, 363, 364, 1010  
 <window frame start> • 356, **357**, 358  
 <window frame units> • **356**  
 <window function> • 61, 62, 188, 189, **207**, 208, 209, 210, 230, 252, 336, 344, 354, 358, 359, 369, 372, 549, 550, 551, 1220, 1221  
 <window function type> • **207**, 209, 372  
 <window name> • **164**, 170, 171, 207, 208, 209, 210, 356, 358, 1221  
 <window name or specification> • **207**, 209, 210, 1221  
 <window order clause> • 62, **356**, 357, 358, 360, 363  
 window ordering • 61, 62, 210, 358, 360, 1221  
 window partition • 61, 63, 64, 360, 1248  
 <window partition clause> • 62, **356**, 357, 358, 360, 481  
 <window partition column reference> • **356**, 357, 359  
 <window partition column reference list> • **356**  
 window partitioning • 61, 62  
 <window specification> • 62, 207, 209, 252, **356**, 365, 1010, 1220  
 <window specification details> • 208, **356**  
 window structure descriptor • 62, 64, 209, 549, 1248  
 windowed table • 61, 62, 63, 64  
 WITH • 10, 15, 34, 35, 37, 56, 60, 76, 94, 100, 102, 120, 121, 122, 148, 158, 170, 175, 179, 180, 183, 224, 225, 226, 227, 246, 259, 274, 288, 291, 296, 326, 327, 328, 329, 332, 334, 338, 370, 371, 376, 391, 477, 556, 574, 576, 582, 625, 626, 627, 631, 632, 633, 668, 675, 686, 687, 688, 745, 749, 755, 766, 768, 771, 772, 773, 774, 775, 776, 777, 778, 781, 782, 783, 786, 789, 790, 791, 792, 793, 795, 796, 799, 800, 851, 869, 877, 880, 881, 882, 884, 888, 889, 890, 903, 911, 924, 931, 993, 1001, 1002, 1009, 1011, 1017, 1018, 1019, 1137, 1181, 1193, 1209, 1210, 1217, 1241  
 with check option violation • 924, 931, 1137, 1151  
 <with clause> • 79, **376**, 377, 387, 391, 1210  
 <with column list> • **376**, 377, 387, 393  
 <with list> • 79, **376**, 377, 378, 387  
 <with list element> • 56, 58, 330, **376**, 377, 378, 380, 381, 387, 393, 396, 563  
 <with or without data> • **574**  
 <with or without time zone> • **175**, 179, 184, 1173, 1268, 1269  
 with-return • 93, 94, 102  
 with-return cursor • 10, 93, 94, 103  
 WITHIN • 148, 548  
 <within group specification> • 205, **548**, 550  
 within-group-varying column reference • 203  
 WITHOUT • 15, 34, 35, 37, 94, 100, 148, 157, 158, 175, 179, 183, 224, 225, 226, 227, 288, 291, 292, 296, 851, 852, 902, 907, 1011, 1017, 1038, 1219

WORK • 147, 955, 957  
 working table • 388  
 WRITE • 147, 948, 949, 970

## — X —

X • 153  
 XQuery captured substring • **10**  
 <XQuery option flag> • **10**, **256**, 259, 265, 267, 273, 274, 276, 277, 281, 284, 287, 426, 427, 1011, 1185  
 <XQuery pattern> • **256**, 258, 259, 265, 267, 273, 274, 281, 284, 287, 426, 427, 1011, 1185  
 XQuery regular expression • **10**  
 XQuery regular expression match • 494, 495  
 XQuery regular expression parenthesized sub-expression • **11**  
 <XQuery replacement string> • **11**, **274**, 277, 278, 284, 287, 1011, 1185, 1233

## — Y —

YEAR • 35, 37, 148, 158, 179, 181, 361, 464, 470, 509, 510, 1002, 1006  
 year-month • 34, 37, 38, 183, 294, 361, 470, 511, 613  
 year-month interval • 179  
 <year-month literal> • **154**, 158, 511  
 <years value> • **154**, 158, 511

## — Z —

zero-length character string • 194, 271, 287, 1146, 1231, 1233  
 ZONE • 15, 34, 35, 37, 147, 157, 158, 175, 179, 180, 183, 224, 225, 226, 227, 259, 288, 289, 291, 292, 296, 338, 370, 371, 477, 973, 1001, 1002, 1009, 1011

