INTERNATIONAL STANDARD

ISO/IEC 9075-1

Third edition
2008-07-15

Information technology — Database languages — SQL —

Part 1: Framework (SQL/Framework)

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

Partie 1: Charpente (SQL/Framework)

Reference number
ISO/IEC 9075-1:2008(E)

--`,,```,,,,````-`-`,,`,,`,`,,`---

© ISO/IEC 2008

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS

Not for Resale

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

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

© ISO/IEC 2008 – All rights reserved

Contents

Page

Foreword. ix
Introduction. x
1 Scope. 1
2 Normative references. 3
2.1 ISO and IEC standards. 3
3 Definitions and use of terms. 5
3.1 Definitions. 5
3.1.1 Definitions provided in this standard. 5
3.2 Use of terms. 6
3.3 Informative elements. 7
4 Concepts. 9
4.1 Caveat. 9
4.2 SQL-environments and their components. 9
4.2.1 SQL-environments. 9
4.2.2 SQL-agents. 9
4.2.3 SQL-implementations. 10
4.2.3.1 SQL-clients. 10
4.2.3.2 SQL-servers. 10
4.2.4 SQL-client modules. 10
4.2.5 User identifiers. 11
4.2.6 Roles. 11
4.2.7 User mapping concepts. 11
4.2.8 Routine mapping concepts. 11
4.2.9 Catalogs and schemas. 11
4.2.9.1 Catalogs. 11
4.2.9.2 SQL-schemas. 12
4.2.9.3 The Information Schema. 12
4.2.9.4 The Definition Schema. 12
4.2.10 Foreign servers and descriptors. 12
4.2.11 Foreign-data wrappers and descriptors. 13
4.2.12 SQL-data. 13
4.3 Tables. 13
4.4 SQL data types. 14
4.4.1 General data type information. 14
4.4.2 The null value. 14

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.4.3 Predefined types. . . . 15
4.4.3.1 Numeric types. . . . 15
4.4.3.2 Character string types. . . . 15
4.4.3.3 Binary string types. . . . 15
4.4.3.4 Boolean type. . . . 15
4.4.3.5 Datetime types. . . . 16
4.4.3.6 Interval types. . . . 16
4.4.3.7 XML type. . . . 16
4.4.4 Constructed atomic types. . . . 16
4.4.4.1 Reference types. . . . 16
4.4.5 Constructed composite types. . . . 17
4.4.5.1 Collection types. . . . 17
4.4.5.2 Row types. . . . 17
4.4.5.3 Fields. . . . 17
4.5 Sites and operations on sites. . . . 17
4.5.1 Sites. . . . 17
4.5.2 Assignment. . . . 18
4.5.3 Nullability. . . . 18
4.6 SQL-schema objects. . . . 18
4.6.1 General SQL-schema object information. . . . 18
4.6.2 Descriptors relating to character sets. . . . 19
4.6.2.1 Character sets. . . . 19
4.6.2.2 Collations. . . . 19
4.6.2.3 Transliterations. . . . 19
4.6.3 Domains and their components. . . . 20
4.6.3.1 Domains. . . . 20
4.6.3.2 Domain constraints. . . . 20
4.6.4 User-defined types. . . . 20
4.6.4.1 Structured types. . . . 20
4.6.4.2 Distinct types. . . . 21
4.6.5 Base tables and their components. . . . 21
4.6.5.1 Base tables. . . . 21
4.6.5.2 Columns. . . . 21
4.6.5.3 Table constraints. . . . 21
4.6.5.4 Triggers. . . . 22
4.6.6 View definitions. . . . 22
4.6.7 Assertions. . . . 22
4.6.8 SQL-server modules (defined in [ISO9075-4]). . . . 22
4.6.9 Schema routines. . . . 22
4.6.10 Sequence generators. . . . 23
4.6.11 Privileges. . . . 23
4.7 Integrity constraints and constraint checking. . . . 23
4.7.1 Constraint checking. . . . 23
4.7.2 Determinism and constraints. . . . 24

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.8 Communication between an SQL-agent and an SQL-server. . . . 24
4.8.1 Host languages. . . . 24
4.8.2 Parameter passing and data type correspondences. . . . 25
4.8.2.1 General parameter passing and data type correspondence information. . . . 25
4.8.2.2 Data type correspondences. . . . 25
4.8.2.3 Locators. . . . 25
4.8.2.4 Status parameters. . . . 25
4.8.2.5 Indicator parameters. . . . 26
4.8.3 Descriptor areas. . . . 26
4.8.4 Diagnostic information. . . . 26
4.8.5 SQL-transactions. . . . 26
4.9 Modules. . . . 27
4.10 Routines. . . . 27
4.10.1 General routine information. . . . 27
4.10.2 Type preserving functions. . . . 28
4.11 SQL-statements. . . . 29
4.11.1 Classes of SQL-statements. . . . 29
4.11.2 SQL-statements classified by function. . . . 29
5 The parts of ISO/IEC 9075. . . . 31
5.1 Overview. . . . 31
5.2 ISO/IEC 9075-1: Framework (SQL/Framework). . . . 31
5.3 ISO/IEC 9075-2: Foundation (SQL/Foundation). . . . 32
5.3.1 Data types specified in [ISO9075-2]. . . . 32
5.3.2 Tables. . . . 32
5.3.3 Bindings methods. . . . 32
5.3.3.1 Embedded SQL. . . . 32
5.3.3.2 Dynamic SQL. . . . 32
5.3.3.3 Direct invocation of SQL. . . . 33
5.3.4 SQL-statements specified in [ISO9075-2]. . . . 33
5.4 ISO/IEC 9075-3: Call Level Interface (SQL/CLI). . . . 34
5.5 ISO/IEC 9075-4: Persistent Stored Modules (SQL/PSM). . . . 34
5.5.1 SQL-statements specified in [ISO9075-4]. . . . 35
5.6 ISO/IEC 9075-9: Management of External Data (SQL/MED). . . . 35
5.7 ISO/IEC 9075-10: Object Language Bindings (SQL/OLB). . . . 35
5.8 ISO/IEC 9075-11: Information and Definition Schemas (SQL/Schemata). . . . 35
5.9 ISO/IEC 9075-13: Java Routines & Types Using the Java Programming Language (SQL/JRT). . . . 36
5.10 ISO/IEC 9075-14: XML-Related Specifications (SQL/XML). . . . 36
6 Notation and conventions used in other parts of ISO/IEC 9075. . . . 37
6.1 Notation taken from [ISO10646]. . . . 37
6.2 Notation provided in this International Standard. . . . 37
6.3 Conventions. . . . 39
6.3.1 Specification of syntactic elements. . . . 39
6.3.2 Specification of the Information and Definition Schemata. . . . 39

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6.3.3 Use of terms. . . . 40
6.3.3.1 Syntactic containment. . . . 40
6.3.3.2 Terms denoting rule requirements. . . . 41
6.3.3.3 Rule evaluation order. . . . 41
6.3.3.4 Conditional rules. . . . 42
6.3.3.5 Syntactic substitution. . . . 42
6.3.3.6 Other terms. . . . 43
6.3.3.7 Exceptions. . . . 43
6.3.3.8 General Rules not terminated on exception conditions. . . . 44
6.3.4 Descriptors. . . . 44
6.3.5 Relationships of parts within ISO/IEC 9075. . . . 45
6.3.5.1 New and modified Clauses, Subclauses, and Annexes. . . . 45
6.3.5.2 New and modified tables and figures. . . . 47
6.3.5.3 Functions. . . . 47
6.3.5.4 New and modified Format items. . . . 47
6.3.5.5 New and modified paragraphs and rules. . . . 48
6.3.5.6 Modified annexes. . . . 49
6.3.6 Subclauses used as subroutines. . . . 49
6.3.7 Index typography. . . . 50
6.3.8 Feature ID and Feature Name. . . . 50
7 Annexes to the parts of ISO/IEC 9075. . . . 53
7.1 SQL conformance summary. . . . 53
7.2 Implementation-defined elements. . . . 53
7.3 Implementation-dependent elements. . . . 53
7.4 Deprecated features. . . . 53
7.5 Incompatibilities with previous versions. . . . 53
7.6 SQL feature taxonomy. . . . 54
8 Conformance. . . . 55
8.1 Minimum conformance. . . . 55
8.2 Conformance to parts. . . . 55
8.3 Conformance to features. . . . 55
8.4 Conformance to SQL packages. . . . 56
8.4.1 Enhanced datetime facilities. . . . 57
8.4.2 Enhanced integrity management. . . . 57
8.4.3 PSM. . . . 57
8.4.4 Basic object support. . . . 58
8.4.5 Enhanced object support. . . . 58
8.4.6 Active database. . . . 58
8.4.7 OLAP. . . . 59
8.4.8 Extensions and options. . . . 59
8.5 SQL flagger. . . . 59
8.6 Claims of conformance. . . . 61
8.6.1 Requirements for SQL applications. . . . 61

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8.6.2 Requirements for SQL-implementations. 61
Annex A Maintenance and interpretation of SQL (informative). **63**
Annex B Implementation-defined elements (informative). **65**
Annex C Implementation-dependent elements (informative). **67**
Annex D Deprecated features (infomative). **69**
Annex E Incompatibilities with ISO/IEC 9075:2003 (informative). **71**
Annex F SQL feature taxonomy (informative). **73**
Annex G Defect reports not addressed in this edition of this part of ISO/IEC 9075 (informative). . **75**
Bibliography. **77**
Index. **79**

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Tables

Table

Page

1 Relationships between externally-invoked and SQL-invoked routines. 28
2 Symbols used in BNF. 37
3 SQL Packages. 56
4 Feature taxonomy for optional features. 73

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

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 technical activity. ISO and IEC technical committees collaborate in 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 9075-1 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 9075-1 cancels and replaces the second edition (ISO/IEC 9075-1:2003), which has been technically revised. It also incorporates Technical Corrigendum ISO/IEC 9075-1:2003/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)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

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 ISO/IEC 9075.

3) Clause 3, "Definitions and use of terms", defines terms used in this and other parts of ISO/IEC 9075.

4) Clause 4, "Concepts", describes the concepts used in ISO/IEC 9075.

5) Clause 5, "The parts of ISO/IEC 9075", summarises the content of each of the parts of ISO/IEC 9075, in terms of the concepts described in Clause 4, "Concepts".

6) Clause 6, "Notation and conventions used in other parts of ISO/IEC 9075", defines notation and conventions used in other parts of ISO/IEC 9075.

7) Clause 7, "Annexes to the parts of ISO/IEC 9075", describes the content of annexes of other parts of ISO/IEC 9075.

8) Clause 8, "Conformance", specifies requirements that apply to claims of conformance to all or some of the parts of ISO/IEC 9075.

9) Annex A, "Maintenance and interpretation of SQL", is an informative Annex. It describes the formal procedures for maintenance and interpretation of ISO/IEC 9075.

10) 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.

11) 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.

12) 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.

13) 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.

14) 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.

15) Annex G, "Defect reports not addressed in this edition of this part 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 this part 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. Any resulting blank space is not significant.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Information technology — Database languages — SQL —

Part 1:
Framework (SQL/Framework)

1 Scope

This part of ISO/IEC 9075 describes the conceptual framework used in other parts of ISO/IEC 9075 to specify the grammar of SQL and the result of processing statements in that language by an SQL-implementation.

This part of ISO/IEC 9075 also defines terms and notation used in the other parts of ISO/IEC 9075.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

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

[ISO9075-2] ISO/IEC 9075-2:2008, *Information technology — Database languages — SQL — Part 2: Foundation (SQL/Foundation)*

[ISO9075-3] ISO/IEC 9075-3:2008, *Information technology — Database languages — SQL — Part 3: Call-Level Interface (SQL/CLI)*

[ISO9075-4] ISO/IEC 9075-4:2008, *Information technology — Database languages — SQL — Part 4: Persistent Stored Modules (SQL/PSM)*

[ISO9075-9] ISO/IEC 9075-9:2008, *Information technology — Database languages — SQL — Part 9: Management of External Data (SQL/MED)*

[ISO9075-10] ISO/IEC 9075-10:2008, *Information technology — Database languages — SQL — Part 10: Object Language Bindings (SQL/OLB)*

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

[ISO9075-13] ISO/IEC 9075-13:2008, *Information technology — Database languages — SQL — Part 13: SQL Routines & Types Using the Java™ Programming Language (SQL/JRT)*

[ISO9075-14] ISO/IEC 9075-14:2008, *Information technology — Database languages — SQL — Part 14: XML-Related Specifications (SQL/XML)*

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

© ISO/IEC 2008 – All rights reserved

--`,,```,,,,````-`-`,,`,,`,`,,`---
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

3 Definitions and use of terms

3.1 Definitions

3.1.1 Definitions provided in this standard

In this document, the definition of a verb defines every voice, mood, and tense of that verb.

For the purposes of this document and other parts of ISO/IEC 9075, the following definitions apply:

3.1.1.1 **atomic**: Incapable of being subdivided.

3.1.1.2 **compilation unit**: A segment of executable code, possibly consisting of one or more subprograms.

3.1.1.3 **data type**: A set of representable values.

3.1.1.4 **descriptor**: A coded description of an SQL object. It includes all of the information about the object that a conforming SQL-implementation requires.

3.1.1.5 **fully qualified (of a name of some SQL object)**: with all optional components specified explicitly.

NOTE 1 — A fully qualified name does not necessarily identify an object uniquely. For example, although a fully qualified specific name, consisting of a catalog name, a schema name and a specific name, uniquely identifies a routine, a fully qualified routine name doesn't necessarily do so.

3.1.1.6 **identifier**: A means by which something is identified.

3.1.1.7 **identify**: To reference something without ambiguity.

3.1.1.8 **implementation-defined**: Possibly differing between SQL-implementations, but specified by the implementor for each particular SQL-implementation.

3.1.1.9 **implementation-dependent**: Possibly differing between SQL-implementations, but not specified by ISO/IEC 9075, and not required to be specified by the implementor for any particular SQL-implementation.

3.1.1.10 **instance (of a value)**: A physical representation of a value. Each instance is at exactly one site. An instance has a data type that is the data type of its value.

3.1.1.11 **null value**: A special value that is used to indicate the absence of any data value.

3.1.1.12 **object (as in 'x object')**: Any *thing*. An *X* object is a component of, or is otherwise associated with, some *X*, and cannot exist independently of that *X*. For example, an SQL object is an object that exists only in the context of SQL; an SQL-schema object is an object that exists in some SQL-schema.

3.1.1.13 **persistent**: Continuing to exist indefinitely, until destroyed deliberately. Referential and cascaded actions are regarded as deliberate. Actions incidental to the termination of an SQL-transaction or an SQL-session are not regarded as deliberate.

© ISO/IEC 2008 – All rights reserved

--`,,```,,,,````-`-`,,`,,`,`,,`---
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

3.1.1.14 **property (of an object)**: An attribute, quality, or characteristic of the object.

3.1.1.15 **row**: A sequence of (field name, value) pairs, the data type of each value being specified by the row type.

3.1.1.16 **scope (of a standard)**: The clause in the standard that defines the subject of the standard and the aspects covered, thereby indicating the limits of applicability of the standard or of particular parts of it.

3.1.1.17 **scope (of a name or declaration)**: The one or more BNF non-terminal symbols within which the name or declaration is effective.

3.1.1.18 **scope (of a reference type)**: The table that a value of that reference type references.

3.1.1.19 **sequence**: An ordered collection of objects that are not necessarily distinct.

3.1.1.20 **site**: A place occupied by an instance of a value of some specified data type (or subtype of that data type).

3.1.1.21 **sizing item**: The value of an implementation-defined item for the SQL-implementation or for a profile.

3.1.1.22 **SQL-connection**: An association between an SQL-client and an SQL-server.

3.1.1.23 **SQL-environment**: The context in which SQL-data exists and SQL-statements are executed.

3.1.1.24 **SQL-implementation**: A processor that processes SQL-statements. A *conforming SQL-implementation* is an SQL-implementation that satisfies the requirements for SQL-implementations as defined in Clause 8, “Conformance”.

3.1.1.25 **SQL-session**: The context within which a single user, from a single SQL-agent, executes a sequence of consecutive SQL-statements over a single SQL-connection.

3.1.1.26 **SQL-statement**: A string of characters that conforms to the Format and Syntax Rules specified in the parts of ISO/IEC 9075.

3.1.1.27 **table**: A collection of rows having an ordered collection of one or more columns. Each column has a name and a data type. Each row has, for each column, exactly one value in the data type of that column.

3.2 Use of terms

The concepts on which ISO/IEC 9075 is based are described in terms of objects, in the usual sense of the word.

An object might be considered to be a component of the object on which it depends. If an objects ceases to exist, then every object dependent on that objects also ceases to exist. The representation of an object is known as a descriptor. The descriptor of an object represents everything that needs to be known about the object. See also Subclause 6.3.4, “Descriptors”.

© ISO/IEC 2008 – All rights reserved
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

3.3 Informative elements

In several places in the body of ISO/IEC 9075, informative notes appear. For example:

NOTE 2 — This is an example of a note.

Those notes do not belong to the normative part ISO/IEC 9075. A claim of conformance to material in a Note is not meaningful.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

--`,,```,,,,````-`-`,,`,,`,`,,`---

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS

Not for Resale

4 Concepts

4.1 Caveat

This Clause describes concepts that are, for the most part, specified precisely in other parts of ISO/IEC 9075. In any case of discrepancy, the specification in the other part is to be presumed correct.

4.2 SQL-environments and their components

4.2.1 SQL-environments

An SQL-environment comprises:

— One SQL-agent.

— One SQL-implementation.

— Zero or more SQL-client modules, containing externally-invoked procedures available to the SQL-agent.

— Zero or more authorization identifiers.

— Zero or more user mappings.

— Zero or more routine mappings.

— Zero or more catalogs, each of which contains one or more SQL-schemas, zero or more foreign server descriptors, and zero or more foreign-data wrapper descriptors.

— The sites, principally base tables, that contain SQL-data, as described by the contents of the schemas. This data may be thought of as “the database”, but the term is not used in ISO/IEC 9075, because it has different meanings in the general context.

4.2.2 SQL-agents

An *SQL-agent* is that which causes the execution of SQL-statements. In the case of the direct invocation of SQL (see Subclause 5.3.3.3, “Direct invocation of SQL”), it is implementation-defined. Alternatively, it may consist of one or more compilation units that, when executed, invoke externally-invoked procedures in an SQL-client module.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.2.3 SQL-implementations

An *SQL-implementation* is a processor that executes SQL-statements, as required by the SQL-agent. An SQL-implementation, as perceived by the SQL-agent, includes one SQL-client, to which that SQL-agent is bound, and one or more SQL-servers. An SQL-implementation can conform to ISO/IEC 9075 without allowing more than one SQL-server to exist in an SQL-environment.

Because an SQL-implementation can be specified only in terms of how it executes SQL-statements, the concept denotes an installed instance of some software (database management system). ISO/IEC 9075 does not distinguish between features of the SQL-implementation that are determined by the software vendor and those determined by the installer.

ISO/IEC 9075 recognizes that SQL-client and SQL-server software may have been obtained from different vendors; it does not specify the method of communication between SQL-client and SQL-server.

4.2.3.1 SQL-clients

An *SQL-client* is a processor, perceived by the SQL-agent as part of the SQL-implementation, that establishes SQL-connections between itself and SQL-servers and maintains a diagnostics area and other state data relating to interactions between itself, the SQL-agent, and the SQL-servers.

When one or more SQL-connections have been established for an SQL-client, the diagnostics area maintained by that SQL-client is a copy of the first diagnostics area in the diagnostics area stack (see Subclause 4.8.4, "Diagnostic information") maintained by the SQL-server of the current SQL-connection. When no SQL-connection exists, the diagnostics area is either empty or contains diagnostics information pertaining to some failed SQL-connection.

4.2.3.2 SQL-servers

Each *SQL-server* is a processor, perceived by the SQL-agent as part of the SQL-implementation, that manages SQL-data.

Each SQL-server:

— Manages the SQL-session taking place over the SQL-connection between itself and the SQL-client.

— Executes SQL-statements received from the SQL-client, receiving and sending data as required.

— Maintains the state of the SQL-session, including the authorization identifier and certain session defaults.

4.2.4 SQL-client modules

An *SQL-client module* is a module that is explicitly created and dropped by implementation-defined mechanisms.

An SQL-client module does not necessarily have a name; if it does, the permitted names are implementation-defined.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

An SQL-client module contains zero or more *externally-invoked procedures*. An externally-invoked procedure consists of a single SQL-statement. An externally-invoked procedure is invoked from a compilation unit of a host language.

Exactly one SQL-client module is associated with an SQL-agent at any time. However, in the case of either direct binding style or SQL/CLI, this may a default SQL-client module whose existence is not apparent to the user.

4.2.5 User identifiers

A *user identifier* represents a user. The means of creating and destroying user identifiers, and their mapping to real users, is not specified by ISO/IEC 9075.

4.2.6 Roles

A *role* is a potential grantee and grantor of privileges and of other roles. A role can also own schemas and other objects.

A *role authorization* permits a grantee (see Subclause 4.6.11, "Privileges") to use every privilege granted to the role. It also indicates whether the role authorization is grantable, in which case the grantee is authorized to grant the role, to revoke a grant of the role, and to destroy the role.

4.2.7 User mapping concepts

A user mapping pairs an authorization identifier with a foreign server descriptor.

4.2.8 Routine mapping concepts

A routine mapping pairs an SQL-invoked routine with a foreign server descriptor.

NOTE 3 — In this revision of ISO/IEC 9075, the SQL-invoked routine that is paired with a foreign server descriptor is restricted to be an SQL-invoked regular function. Possible future revisions may drop this restriction.

4.2.9 Catalogs and schemas

4.2.9.1 Catalogs

A *catalog* is a named collection of SQL-schemas, foreign server descriptors, and foreign data wrapper descriptors in an SQL-environment. The mechanisms for creating and destroying catalogs are implementation-defined.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

The default catalog name for <preparable statement>s that are dynamically prepared in the current SQL-session through the execution of <prepare statement>s and <execute immediate statement>s is initially implementation-defined but may be changed by the use of <set catalog statement>s.

4.2.9.2 SQL-schemas

An *SQL-schema*, often referred to simply as a *schema*, is a persistent, named collection of descriptors. Any object whose descriptor is in some SQL-schema is known as an *SQL-schema object*.

A schema, the schema objects in it, and the SQL-data described by them are said to be owned by the authorization identifier associated with the schema.

SQL-schemas are created and destroyed by execution of SQL-schema statements (or by implementation-defined mechanisms).

4.2.9.3 The Information Schema

Every catalog contains an SQL-schema with the name INFORMATION_SCHEMA that includes the descriptors of a number of schema objects, mostly view definitions, that together allow every descriptor in that catalog to accessed, but not changed, as though it was SQL-data.

The data available through the views in an Information Schema includes the descriptors of the Information Schema itself. It does not include the schema objects or base tables of the Definition Schema (see Subclause 4.2.9.4, "The Definition Schema").

Each Information Schema view is so specified that a given user can access only those rows of the view that represent descriptors on which that user has privileges.

4.2.9.4 The Definition Schema

The *definition schema* is a fictitious schema with the name DEFINITION_SCHEMA; if it were to exist, the SQL-data in its base tables would describe all the SQL-schemas available to an SQL-server. ISO/IEC 9075 defines it only in order to use it as the basis for the views of the Information Schema.

The structure of the Definition Schema is a representation of the data model of SQL.

4.2.10 Foreign servers and descriptors

A foreign server is a processor that is not part of the SQL-implementation. A foreign server is described by a foreign server descriptor. A foreign server manages data that is not part of the SQL-environment. An SQL-server and an SQL-client can use a foreign server descriptor, which is a catalog element, to communicate with a foreign server. The data managed by a foreign server can be accessed by an SQL-server or an SQL-client through foreign tables, which are SQL-schema elements.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.2.11 Foreign-data wrappers and descriptors

A foreign-data wrapper provides the mechanism by which an SQL-server can access the data that is managed by a foreign server. A foreign-data wrapper is described by a foreign-data wrapper descriptor.

4.2.12 SQL-data

SQL-data is data described by SQL-schemas — data that is under the control of an SQL-implementation in an SQL-environment.

4.3 Tables

A *table* has an ordered collection of one or more columns and a collection of zero or more rows. In general, rows in a table are unordered; however, rows in a table are ordered if the table is the result of a <query expression> that immediately contains an <order by clause>. Each column has a name and a data type. Each row has exactly one value for each column *C*, and the data type of that value is the data type of *C*.

SQL-data consists entirely of table variables, called *base tables*. An operation that references zero or more base tables and returns a table is called a *query*. The result of a query is called a *derived table*. The rows of a derived table may be ordered.

The rows of a table have a type, called “the row type”; every row of a table has the same row type, which is also the row type of the table. A table that is declared to be based on some structured type is called a “typed table”; its columns correspond in name and declared type to the attributes of the structured type. Typed tables have one additional column, called the “self-referencing column” whose type is a reference type associated with the structured type of the table.

If a typed table *TB1* has an associated structured type *TP1* that is a subtype of some other structured type *TP2*, then *TB1* can be defined to be a “subtable” of a typed table *TB2* whose associated type is *TP2*; *TB2* is, in this case, a “supertable” of *TB1*.

A *view* is a named query, which can be invoked by use of this name. The result of such an invocation is called a *viewed table*.

Some queries, and hence some views, are *updatable*, meaning they can appear as targets of statements that change SQL-data. The results of changes expressed in this way are defined in terms of corresponding changes to base tables.

No two columns of a base table or a viewed table can have the same name. Derived tables, other than viewed tables, may contain more than one column with the same name.

A base table is either a schema object (its descriptor is in a schema; see Subclause 4.6.5, “Base tables and their components”) or a module object (its descriptor is in a module; see Subclause 4.9, “Modules”). A base table whose descriptor is in a schema is called a *created base table*, and may be either persistent or temporary (though its descriptor is persistent in either case). A *persistent base table* contains 0 (zero) or more rows of persistent SQL-data. A base table declared in a module may only temporary, and is called a *declared temporary table*.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

A *temporary table* is an SQL-session object that cannot accessed from any other SQL-session. A *global temporary table* can be accessed from any associated SQL-client module. A *local temporary table* can be accessed only from the module to which it is local.

A temporary table is empty when an SQL-session is initiated and it is emptied (that is, all its rows are deleted) either when an SQL-transaction is terminated or when an SQL-session is terminated, depending on its descriptor.

4.4 SQL data types

4.4.1 General data type information

Every data value belongs to some data type.

Every data type is either *predefined*, *constructed*, or *user-defined*. Every data type has a name. The name of a predefined or constructed data type is a reserved word specified by that part of ISO/IEC 9075 that specifies the data type. The name of a user-defined type is provided in its definition. A user-defined data type is a schema object; see Subclause 4.6.4, "User-defined types".

A predefined data type is a data type specified by ISO/IEC 9075, and is therefore provided by the SQL-implementation. A data type is predefined even though the user is required (or allowed) to provide certain parameters when specifying it (for example the precision of a number).

A predefined data type is atomic. An atomic type is a data type whose values are not composed of values of other data types. The existence of an operation (SUBSTRING, EXTRACT) that is capable of selecting part of a string or datetime value does not imply that a string or datetime is not atomic.

A constructed type is either atomic or composite. A composite type is a data type each of whose values is composed of zero or more values, each of a declared data type.

4.4.2 The null value

Every data type includes a special value, called the *null value*, sometimes denoted by the keyword NULL. This value differs from other values in the following respects:

— Since the null value is in every data type, the data type of the null value implied by the keyword NULL cannot be inferred; hence NULL can be used to denote the null value only in certain contexts, rather than everywhere that a literal is permitted.

— Although the null value is neither equal to any other value nor not equal to any other value — it is *unknown* whether or not it is equal to any given value — in some contexts, multiple null values are treated together; for example, the <group by clause> treats all null values together.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.4.3 Predefined types

4.4.3.1 Numeric types

There are two classes of numeric type: *exact numeric*, which includes integer types and types with specified precision and scale; and *approximate numeric*, which is essentially floating point, and for which a precision may optionally be specified.

Every number has a *precision* (number of digits), and exact numeric types also have a scale (digits after the radix point). Arithmetic operations may be performed on operands of different or the same numeric type, and the result is of a numeric type that depends only on the numeric type of the operands. If the result cannot be represented exactly in the result type, then whether it is rounded or truncated is implementation-defined. An exception condition is raised if the result is outside the range of numeric values of the result type, or if the arithmetic operation is not defined for the operands.

4.4.3.2 Character string types

A value of *character string type* is a string (sequence) of characters drawn from some character repertoire. The characters in a character string *S* are all drawn from the same character set *CS*. If *S* is the value of some expression *E*, then *CS* is the character set specified for the declared type of *E*. A character string type is either of fixed length, or of variable length up to some implementation-defined maximum. A value of *character large object type* is a string of characters from some character repertoire and is always associated with exactly one character set. A *large object character string* is of variable length, up to some implementation-defined maximum that is probably greater than that of other character strings.

A character string may be specified as being based on a specific character set by specifying CHARACTER SET in the data type; a particular character set chosen by the implementation to be the *national character set* may be specified by specifying NATIONAL CHARACTER, NATIONAL CHARACTER VARYING, or NATIONAL CHARACTER LARGE OBJECT (or one of several syntactic equivalents) as the data type.

4.4.3.3 Binary string types

A value of *binary string type* is a string (sequence) of octets. A binary string type is either of fixed length or of variable length, up to an implementation-defined maximum. A value of *binary large object type* is a string of octets. A *binary large object string* is of variable length, up to some implementation-defined maximum that is probably greater than that of other binary strings.

4.4.3.4 Boolean type

A value of the *Boolean* data type is either *true* or *false*. The truth value of *unknown* is sometimes represented by the null value.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.4.3.5 Datetime types

There are three *datetime types*, each of which specifies values comprising datetime fields.

A value of data type TIMESTAMP comprises values of the datetime fields YEAR (between 0001 and 9999), MONTH, DAY, HOUR, MINUTE and SECOND.

A value of data type TIME comprises values of the datetime fields HOUR, MINUTE and SECOND.

A value of data type DATE comprises values of the datetime fields YEAR (between 0001 and 9999), MONTH, and DAY.

A value of DATE is a valid Gregorian date. A value of TIME is a valid time of day.

TIMESTAMP and TIME may be specified with a number of (decimal) digits of fractional seconds precision.

TIMESTAMP and TIME may also be specified as being WITH TIME ZONE, in which case every value has associated with it a time zone displacement. In comparing values of a data type WITH TIME ZONE, the value of the time zone displacement is disregarded.

4.4.3.6 Interval types

A value of an *interval type* represents the duration of a period of time. There are two classes of intervals. One class, called *year-month intervals*, has a datetime precision that includes a YEAR field or a MONTH field, or both. The other class, called *day-time intervals*, has an express or implied interval precision that can include any set of contiguous fields other than YEAR or MONTH.

4.4.3.7 XML type

Values of the XML type are called XML values.

4.4.4 Constructed atomic types

4.4.4.1 Reference types

A *reference type* is a constructed data type, a value of which references (or points to) some site holding a value of the referenced type. The only sites that may be so referenced are the rows of typed tables. It follows that every referenced type is a structured type.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.4.5 Constructed composite types

4.4.5.1 Collection types

A *collection* comprises zero or more elements of a specified data type known as the *element type*.

An *array* is an ordered collection of not necessarily distinct values, whose elements are referenced by their ordinal position in the array.

An array type is specified by an array type constructor.

A multiset is an unordered collection of not necessarily distinct values.

A multiset type is specified by a multiset type constructor.

4.4.5.2 Row types

A row type is a sequence of one or more (field name, data type) pairs, known as fields. A value of a row type consists of one value for each of its fields.

4.4.5.3 Fields

A field is a (field name, data type) pair. A value of a field is a value of its data type.

4.5 Sites and operations on sites

4.5.1 Sites

A *site* is a place that holds an instance of a value of a specified data type. Every site has a defined degree of persistence, independent of its data type. A site that exists until deliberately destroyed is said to be *persistent*. A site that necessarily ceases to exist on completion of a compound SQL-statement, at the end of an SQL-transaction, or at the end of an SQL-session is said to *temporary*. A site that exists only for as long as necessary to hold an argument or returned value is said to be *transient*.

As indicated above, the principal kind of persistent or temporary site is the base table. A base table is a special kind of site, in that constraints can be specified on its values, which the SQL-implementation is required to enforce (see Subclause 4.6.5.3, “Table constraints”).

Some sites may be referenced by their names — for example, base tables and SQL variables (see [ISO9075-4]). Some sites may be referenced by a REF value. A site occupied by an element of an array may referenced by its element number.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.5.2 Assignment

The instance at a site can be changed by the operation of *assignment*. Assignment replaces the instance at a site (known as the *target*) with a new instance of a (possibly different) value (known as the *source* value). Assignment has no effect on the reference value of a site, if any.

4.5.3 Nullability

Every site has a *nullability characteristic*, which indicates whether it may contain the null value (is *possibly nullable*) or not (is *known not nullable*). Only the columns of base tables may be constrained to be known not nullable, but columns derived from such columns may inherit the characteristic.

No table can be null, though a table may have no rows.

4.6 SQL-schema objects

4.6.1 General SQL-schema object information

An SQL-schema object has a descriptor. The descriptor of a persistent base table describes a persistent object that has a separate, though dependent, existence as SQL-data. Other descriptors describe SQL objects that have no existence distinct from their descriptors (at least as far as ISO/IEC 9075 is concerned). Hence there is no loss of precision if, for example, the term “assertion” is used when “assertion descriptor” would be more strictly correct.

Every schema object has a name that is unique within the schema among objects of the name class to which it belongs. The name classes are:

— Base tables and views.

— Domains and user-defined types.

— Table constraints, domain constraints, and assertions.

— SQL-server modules.

— Triggers.

— SQL-invoked routines (specific names only, which are not required to be specified explicitly, but if not are implementation-dependent).

— Character sets.

— Collations.

— Transliterations.

— Sequence generators.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Certain schema objects have named components whose names are required to be unique within the object to which they belong. Thus columns are uniquely named components of base tables or views, attributes are uniquely named components of structured types, and fields are uniquely named components of row types.

Some schema objects may be provided by the SQL-implementation and can neither created nor dropped by a user.

4.6.2 Descriptors relating to character sets

4.6.2.1 Character sets

A *character set* has a named set of characters (*character repertoire*) that may be used for forming values of the character string type, as well as a named *character encoding form*. Every character set has a *default collation*. Character sets provided by the SQL-implementation, whether defined by other standards or by the SQL-implementation, are represented in the Information Schema.

When characters are contained entirely within an SQL-implementation, the methods for encoding them and for collecting them into strings are implementation-dependent. When characters are exchanged with host programs or other entities outside of the SQL-implementation, the character sets also have an encoding, which specifies the bits used to represent each character, and a *character encoding form*, which specifies the scheme used to collect characters into strings.

The character repertoire of every character set supported by an SQL-implementation is some subset of the repertoire of the Universal Character Set specified by [UCS].

The notation specified in [UCS], Subclause 6.5, "Short identifiers for code positions (UIDs)", is the canonical representation of characters and character strings in ISO/IEC 9075.

NOTE 4 — [UCS] assigns a range of code points for “private use”. Future editions of [UCS] are likely to add more code points, which SQL-implementations are required to support.

4.6.2.2 Collations

A *collation* is a named operation for ordering character strings in a particular character repertoire.

A site declared with a character string type may be specified as having a collation, which is treated as part of its data type.

Every collation shall be defined by or derived from an International Standard such as ISO/IEC 14561 or by a National Standard, or shall be an implementation-defined collation.

4.6.2.3 Transliterations

A *transliteration* is a named operation for mapping from a character string of some character set into a character string of a given, not necessarily distinct, character set. The operation is performed by invocation of an external function identified by the name of the transliteration. Since an entire string is passed to this function and a string

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

returned, the mapping is not necessarily from one character to one character, but may be from a sequence of one or more characters to another sequence of one or more characters.

4.6.3 Domains and their components

4.6.3.1 Domains

A *domain* is a named user-defined object that can be specified as an alternative to a data type in certain places where a data type can be specified. A domain consists of a data type, possibly a default option, and zero or more (domain) constraints.

4.6.3.2 Domain constraints

A *domain constraint* applies to every column that is based on that domain, by operating as a table constraint for each such column.

Domain constraints apply only to columns based on the associated domain.

A domain constraint is applied to any value resulting from a cast operation to the domain.

4.6.4 User-defined types

4.6.4.1 Structured types

A *structured type* is a named, user-defined data type. A value of a structured type comprises a number of *attribute values*. Each attribute of a structured type has a data type, specified by an *attribute type* that is included in the descriptor of the structured type.

Attribute values are said to be *encapsulated*; that is to say, they are not directly accessible to the user, if at all. An attribute value is accessible only by invoking a function known as an *observer function* that returns that value. A value of a structured type can also be accessed by a *locator*.

A structured type may be defined to be a *subtype* of another structured type, known as its *direct supertype*. A subtype *inherits* every attribute of its direct supertype, and may have additional attributes of its own. An expression of a subtype may appear anywhere that an expression of any of its supertypes is allowed (this concept is known as *substitutability*). Moreover, the value of an expression may be a value of any subtype of the declared type of the expression.

One or more base tables can be created based on a structured type. A base table based on a structured type *ST* can be a *subtable* of a base table based on a supertype of *ST*.

One or more views can be created based on a structured type. A view based on a structured type *ST* can be a subview of a view based on a supertype of *ST*.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

An *attribute* is a named component of a structured type. It has a data type and a default value.

4.6.4.2 Distinct types

A *distinct type* is a user-defined data type that is based on some predefined type. The values of a distinct type are represented by the values of the type on which it is based.

An argument of a distinct type can be passed only to a parameter of the same distinct type. This allows precise control of what routines can invoked on arguments of that data type.

4.6.5 Base tables and their components

4.6.5.1 Base tables

A *base table* is a site that holds a table value (see Subclause 4.3, "Tables"). All SQL-data is held in base tables.

If a base table is based on a structured type, it may be a *subtable* of one or more other base tables that are its *supertables*.

4.6.5.2 Columns

A *column* is a named component of a table. It has a data type, a default, and a nullability characteristic.

4.6.5.3 Table constraints

A *table constraint* is an integrity constraint associated with a single base table.

A table constraint is either a *unique constraint*, a *primary key constraint*, a *referential constraint*, or a *check constraint*.

A unique constraint specifies one or more columns of the table as *unique columns*. 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.

A primary key constraint is a unique constraint that specifies PRIMARY KEY. A primary key constraint is satisfied if and only if no two rows in a table have the same non-null values in the unique columns and none of the values in the specified column or columns are the null value.

A referential constraint specifies one or more columns as *referencing columns* and corresponding *referenced columns* in some (not necessarily distinct) base table, referred to as the *referenced table*. Such referenced columns are the unique columns of some unique constraint of the referenced table. A referential constraint is always satisfied if, for every row in the referencing table, the values of the referencing columns are equal to those of the corresponding referenced columns of some row in the referenced table. If null values are present,

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

however, satisfaction of the referential constraint depends on the treatment specified for nulls (known as the *match type*).

Referential actions may be specified to determine what changes are to be made to the referencing table if a change to the referenced table would otherwise cause the referential constraint to be violated.

A table check constraint specifies a *search condition*. The constraint is violated if the result of the search condition is *False* for any row of the table (but not if it is *Unknown*).

4.6.5.4 Triggers

A *trigger* is an object associated with a single base table or a single viewed table. A trigger specifies a *subject table*, *trigger event*, a *trigger action time*, and one or more *triggered actions*.

A trigger event specifies what action on the subject table shall cause the triggered actions. A trigger event is either INSERT, DELETE, or UPDATE.

A trigger action time specifies whether the triggered action is to taken BEFORE, INSTEAD OF, or AFTER the trigger event.

A triggered action is either an SQL procedure statement or BEGIN ATOMIC, followed by one or more <SQL procedure statement>s terminated with <semicolon>s, followed by END.

4.6.6 View definitions

A *view* (strictly, a *view definition*) is a named query, that may for many purposes be used in the same way as a base table. Its value is the result of evaluating the query. See also Subclause 4.3, “Tables”.

4.6.7 Assertions

An *assertion* is a check constraint. The constraint is violated if the result of the search condition is false (but not if it is unknown).

4.6.8 SQL-server modules (defined in [ISO9075-4])

An *SQL-server module* is a module that is a schema object. See Subclause 4.9, “Modules”.

4.6.9 Schema routines

A *schema routine* is an SQL-invoked routine that is a schema object. See Subclause 4.10, “Routines”.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.6.10 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 sequence generator has a time-varying *current base value*, which is a value of its data type. Refer to Subclause 4.21, "Sequence generators", in [ISO9075-2] for details.

4.6.11 Privileges

A *privilege* represents a grant, by some grantor, to a specified grantee (which is either an authorization identifier, a role, or PUBLIC), of the authority required to use, or to perform a specified action on, a specified schema object. The specifiable actions are: SELECT, INSERT, UPDATE, DELETE, REFERENCES, USAGE, UNDER, TRIGGER, and EXECUTE.

A privilege *with grant option* authorizes the grantee to grant that privilege to other grantees, with or without the grant option.

A SELECT privilege *with hierarchy option* automatically provides the grantee with SELECT privileges on all subtables, both existing and any that may be added in the future, on the table on which the privilege is granted.

Every possible grantee is authorized by privileges granted to PUBLIC. SELECT with grant option is granted to PUBLIC for every schema object in the Information Schema.

An authorization identifier who creates a schema object is automatically granted all possible privileges on it, with grant option.

Only an authorization identifier who has some privilege on a schema object is able to discover its existence.

4.7 Integrity constraints and constraint checking

4.7.1 Constraint checking

There are two kinds of schema object that describe constraints: assertions and table constraints (including domain constraints of any domains on which columns of that table may be based), and they are checked in the same way.

Every constraint is either *deferrable* or *not deferrable*.

In every SQL-session, every constraint has a *constraint mode* that is a property of that SQL-session. Each constraint has a (persistent) default constraint mode, with which the constraint starts each SQL-transaction in each SQL-session.

A constraint mode is either *deferred* or *immediate*, and can be set by an SQL-statement, provided the constraint is deferrable.

When a transaction is initiated, the constraint mode of each constraint is set to its default.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

On completion of execution of every SQL-statement, every constraint is checked whose constraint mode is immediate.

Before termination of a transaction, every constraint mode is set to immediate (and therefore checked).

4.7.2 Determinism and constraints

Expression evaluation results may be non-deterministic. For example, in the case of a column whose data type is varying character string, the value remaining after the elimination of duplicates may be different on different occasions, even though the data is the same. This can occur because the number of trailing spaces may vary from one duplicate to another, and the value to be retained, after the duplicates have been eliminated, is not specified by ISO/IEC 9075. Hence, the length of that value is non-deterministic. In such a case, the expression, and any expression whose value is derived from it, is said to be *possibly non-deterministic* ("possibly", because it may be that all SQL-agents that ever update that column may remove trailing spaces; but this cannot be known to the SQL-implementation).

Because a constraint that contains a possibly non-deterministic expression might be satisfied at one time, yet fail at some later time, no constraint is permitted to contain such an expression.

A routine may claim to be deterministic; if it isn't, then the effect of invoking the routine is implementation-dependent.

4.8 Communication between an SQL-agent and an SQL-server

4.8.1 Host languages

A *host language* is a programming language that can be used to write an SQL-agent. For an SQL-implementation to communicate successfully with an SQL-agent, the latter must be written using a host language supported by the SQL-implementation. The set of host languages supported by an SQL-implementation is implementation-defined. [ISO9075-2], [ISO9075-3], and [ISO9075-10] define bindings between an SQL-implementation and several programming languages, mostly languages for which there is an International Standard. A conforming SQL-implementation is required to support at least one host language for which ISO/IEC 9075 defines a binding.

There are several methods of communicating, known as *binding styles*.

— The SQL-client module binding style (specified in [ISO9075-2]). In this binding style, the user, using an implementation-defined mechanism, specifies a module to be used as an SQL-client module.

— The Call-Level Interface (specified in [ISO9075-3]). In this case, the SQL-agent invokes one of a number of standard routines, passing appropriate arguments, such as a character string whose value is some SQL-statement.

— Embedded SQL (specified in [ISO9075-2] and [ISO9075-10]). In this case, SQL-statements are coded into the application program; an implementation-dependent mechanism is then used to:

  • Generate from each SQL-statement an externally-invoked procedure. These procedures are collected together into a module, for subsequent use as an SQL-client module. In Object Language Bindings, described in [ISO9075-10], the roles of modules and procedures are served by classes and their methods.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

- Replace each SQL-statement with an invocation of the externally-invoked procedure or method generated from it.

— Direct invocation of SQL (specified in [ISO9075-2]). Direct invocation is a method of executing SQL-statements directly, through a front-end that communicates directly with the user.

No matter what binding style is chosen, SQL-statements are written in an implementation-defined character set, known as the *source language character set*. The source language character set is not required to be the same as the character set of any character string appearing in SQL-data.

[ISO9075-2] specifies the actions of an externally-invoked procedure in an SQL-client module when it is called by a host language program. [ISO9075-10] specifies the actions of a method of an SQL-client class when it is invoked from a host program.

4.8.2 Parameter passing and data type correspondences

4.8.2.1 General parameter passing and data type correspondence information

Each parameter in the parameter list of an externally-invoked procedure has a name and a data type. The method and time of binding between an external routine's reference to a compilation unit and that compilation unit is implementation-defined.

4.8.2.2 Data type correspondences

ISO/IEC 9075 specifies correspondences between SQL data types and host language data types. Not every SQL data type has a corresponding data type in every host language.

4.8.2.3 Locators

A host variable, host parameter, SQL parameter of an external routine, or the value returned by an external function may be specified to be a *locator*. The purpose of a locator is to allow very large data instances to be operated upon without transferring their entire value to and from the SQL-agent. Refer to Subclause 4.29.5, “Locators”, in [ISO9075-2] for details.

4.8.2.4 Status parameters

Every externally-invoked procedure is required to have an output parameter called SQLSTATE, which is known as a *status parameter*.

SQLSTATE is a character string of length 5, whose values are defined by the parts of ISO/IEC 9075. An SQLSTATE value of '00000' (five zeros) indicates that the most recent invocation of an externally-invoked procedure was successful.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.8.2.5 Indicator parameters

An *indicator parameter* is an integer parameter that, by being specified immediately following a parameter (other than an indicator parameter), is associated with it. A negative value in an indicator parameter indicates that the associated parameter is null. A value greater than zero indicates what the length of the value of the associated parameter would have been, had it not been necessary to truncate it. This may arise with a character string and with certain other data types.

If a null value is to be assigned to a parameter that has no associated indicator parameter, then an exception condition is raised.

4.8.3 Descriptor areas

A *descriptor area* (not to be confused with a descriptor) is a named area allocated by the SQL-implementation at the request of the SQL-agent. A descriptor area is used for communication between the SQL-implementation and the SQL-agent. There are SQL-statements for transferring information between the SQL-agent and a descriptor area.

4.8.4 Diagnostic information

A *diagnostics area* is a communication area allocated by the SQL-implementation in which one or more *conditions* are recorded as they arise.

Whenever the SQL-implementation executes an SQL-statement (other than an SQL-diagnostics statement), it sets values, representing one or more conditions resulting from that execution, in a diagnostics area. These values give some indication of what has happened. They can be accessed by SQL-diagnostics statements. Execution of an SQL-diagnostics statement does not cause any conditions to be recorded.

A conforming SQL-implementation is not required to set more than one condition at the same time.

Multiple diagnostics areas can exist when an SQL-statement *SS* is being executed by an SQL-server. This happens in certain specific circumstances when execution of *SS* causes other SQL-statements to be executed, synchronously, before the execution of *SS* is complete. The multiple diagnostics areas form a pushdown stack, the first element of which contains information pertaining to the most recently executed SQL-statement. Two examples of when additional diagnostics areas come into existence are when *SS* is or contains a <call statement> and when execution of *SS* causes a triggered action to be executed.

4.8.5 SQL-transactions

An *SQL-transaction* (*transaction*) is a sequence of executions of SQL-statements that is atomic with respect to recovery. That is to say: either the execution result is completely successful, or it has no effect on any SQL-schemas or SQL-data.

At any time, there is at most one current SQL-transaction between the SQL-agent and the SQL-implementation.

If there is no current SQL-transaction, execution of a *transaction-initiating statement* or evaluation of a <subquery> will initiate one.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Every SQL-transaction is terminated by either a commit statement or a rollback statement. The execution of either of these statements may be implicit.

An SQL-transaction has a *transaction state*. Certain properties of the transaction state are set by the execution of SQL-statements. Such SQL-statements may be executed only when there is no SQL-transaction current. On the first occasion, at or after a transaction is initiated, that the SQL-client connects to, or sets the connection to, an SQL-server, the properties are sent to that SQL-server.

The *access mode* of an SQL-transaction indicates whether the transaction is read-only (is not permitted to change any persistent SQL-data) or read-write (is permitted to change persistent SQL-data).

The *isolation level* of an SQL-transaction specifies the extent to which the effects of actions by SQL-agents other than in the SQL-environment, are perceived within that SQL-transaction.

Every isolation level guarantees that every SQL-transaction is executed; SQL-transactions not executing completely fail completely. Every isolation level guarantees that no update is lost. The highest isolation level, SERIALIZABLE, guarantees *serializable* execution, meaning that the effect of SQL-transactions that overlap in time is the same as the effect they would have had, had they not overlapped in time. The other levels of isolation, REPEATABLE READ, READ COMMITTED, and READ UNCOMMITTED, guarantee progressively lower degrees of isolation.

4.9 Modules

There are three kinds of modules, each of which has certain properties and contains various kinds of module objects (also known as module contents). The principal module objects are one or more routines (see Subclause 4.10, “Routines”).

A module is one of the following:

— An SQL-client module, containing only externally invoked procedures.

— An SQL-server module, containing only SQL-invoked routines.

4.10 Routines

4.10.1 General routine information

Table 1, “Relationships between externally-invoked and SQL-invoked routines”, shows the terms used for the various possible combinations of routine written in SQL and written in another language, as well as routines invoked from SQL and routines invoked from another language.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Table 1 — Relationships between externally-invoked and SQL-invoked routines

	Routines written in SQL	**Routines written in languages other than SQL**
Routines invoked from SQL	SQL functions and SQL procedures	External functions and procedures
Routines invoked from languages other than SQL	Externally-invoked procedures	*(not relevant to SQL)*

An *SQL-invoked routine* is a routine that can be invoked from SQL. It is either a function or a procedure. Some functions have special properties that characterize them as *methods*.

An SQL-invoked routine is either a schema object or a component of an SQL-server module (itself a schema object).

An *SQL-invoked procedure* is a procedure invoked by an SQL call statement. An *SQL-invoked function* is invoked by a routine invocation in some value expression.

An *external routine* is an SQL-invoked routine that references some compilation unit of a specified programming language that is outside the SQL-environment. The mechanism and time of binding of such a reference is implementation-defined.

An *SQL routine* is an SQL-invoked routine whose routine body is written in SQL.

The name of an SQL-invoked routine is not required to be unique. If two or more routines share the same name, that name is said to be *overloaded*, and an invocation of that name will cause execution of the routine whose signature best matches the arguments of the invocation. Normally, only the declared types of the expressions denoting the argument values are considered in determining the best match, but in the case of methods, which are invoked using a distinguishing syntax, the most specific type of one of the arguments is taken into consideration.

4.10.2 Type preserving functions

If an SQL-invoked function has a parameter that is specified RESULT, that parameter is known as a *result parameter*, and if the data type of the result parameter and that of the result are the same user-defined type, then the function is said to be a *type preserving function*. The result of a type preserving function is the value of the result parameter, possibly mutated.

Every mutator function is a type preserving function.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

4.11 SQL-statements

4.11.1 Classes of SQL-statements

An SQL-statement is a string of characters that conforms to the Format and Syntax Rules specified in one of the parts of ISO/IEC 9075.

Most SQL-statements can be prepared for execution and executed in an SQL-client module. In this case, an externally-invoked procedure with a single SQL-statement is created when the SQL-statement is prepared and that externally-invoked procedure is implicitly called whenever the prepared SQL-statement is executed.

There are at least five ways of classifying SQL-statements:

— According to their effect on SQL objects, whether persistent objects (*i.e.*, SQL-data, SQL-schemas and their contents, or SQL-client modules) or temporary objects (such as SQL-sessions and other SQL-statements).

— According to whether or not they initiate an SQL-transaction, or can, or shall, be executed when no SQL-transaction is active.

— According to whether or not they may be embedded in a program written in a host language.

— According to whether or not they may be directly executed.

— According to whether or not they may be dynamically prepared and executed.

ISO/IEC 9075 permits implementations to provide additional, implementation-defined, statements that may fall into any of these categories.

4.11.2 SQL-statements classified by function

The following are the broad classes of SQL-statements:

— SQL-schema statements that can be used to create, alter, and drop schemas and schema objects.

— SQL-data statements that perform queries, and insert, update, and delete operations on tables. Execution of an SQL-data statement is capable of affecting more than one row, of more than one table.

— SQL-transaction statements that set parameters for, and start or terminate transactions.

— SQL-control statements that may be used to control the execution of a sequence of SQL statements.

— SQL-connection statements that initiate and terminate connections, and allow an SQL-client to switch from an SQL-session with one SQL-server to an SQL-session with another.

— SQL-session statements that set some default values and other parameters of an SQL-session.

— SQL-diagnostics statements that get diagnostics (from the diagnostics area) and signal exceptions in SQL routines.

— SQL-dynamic statements that support the preparation and execution of dynamically-generated SQL-statements, and obtaining information about them.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

5 The parts of ISO/IEC 9075

5.1 Overview

This part of ISO/IEC 9075, Framework, is a prerequisite for all other parts, because it describes the basic concepts on which other parts are based and the notation used in them.

[ISO9075-2], Foundation, specifies the structure of SQL-statements and the effects of executing them.

Every part of ISO/IEC 9075 other than parts 1 and 2 is specified as an amendment to other parts of ISO/IEC 9075. See Subclause 6.3.5, "Relationships of parts within ISO/IEC 9075", for details.

[ISO9075-3], Call-level interface, specifies another mechanism of communication between an SQL-agent and an SQL-implementation.

[ISO9075-4], Persistent Stored Modules, specifies significant additions to SQL itself by making SQL computationally complete.

[ISO9075-9], Management of External Data, specifies significant additions to SQL that permit an SQL-agent to access data not under the control of SQL-servers in the SQL-environment.

[ISO9075-10], Object Language Bindings, specifies the manner in which SQL statements can be embedded into Java programs.

[ISO9075-11], Information and Definition Schemas, specifies views by which an application may learn the names of persistent database objects, such as tables, views, columns, *etc*.

[ISO9075-13], Routines & Types Using the Java Programming Language, specifies the ability to invoke static methods written in the Java programming language as SQL-invoked routines and to use classes defined in the Java programming language as SQL structured types.

[ISO9075-14], XML-Related Specifications, defines ways in which Database Language SQL can be used in conjunction with XML.

The content of each part is described in the following subclauses.

5.2 ISO/IEC 9075-1: Framework (SQL/Framework)

This part of ISO/IEC 9075 contains:

— A description of an SQL-environment, and brief descriptions of the concepts used in ISO/IEC 9075.

— A brief description of the content of each part. These descriptions are purely informative, and do not constitute requirements.

— Notations and conventions that apply to all or most parts of ISO/IEC 9075. Other parts specify further conventions as required.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

5.3 ISO/IEC 9075-2: Foundation (SQL/Foundation)

[ISO9075-2] specifies the following features of SQL.

5.3.1 Data types specified in [ISO9075-2]

The following data types are specified in [ISO9075-2]:

— All numeric and string types.

— The Boolean type.

— All datetime and interval types.

— Row types.

— Collection types.

— User-defined types.

— Domains.

— Reference types.

5.3.2 Tables

Rules for determining functional dependencies and candidate keys of tables are defined.

5.3.3 Bindings methods

[ISO9075-2] specifies three methods of binding an SQL-agent to an SQL-implementation.

5.3.3.1 Embedded SQL

Embedded SQL is a method of embedding SQL-statements in a compilation unit that is otherwise written according to the rules for a particular programming language, known as the *host language*. It defines how an equivalent compilation unit, entirely in the host language, may be derived. In that equivalent compilation unit, each embedded SQL-statement has been replaced by one or more statements that invoke a database language procedure that contains the SQL-statement.

5.3.3.2 Dynamic SQL

Dynamic SQL is an extension of Embedded SQL. Facilities are specified to:

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

— Allocate and free a descriptor area used for communication between the SQL-implementation and the SQL-agent.

— Cause the execution of SQL-statements, including the preparation of statements for subsequent execution.

5.3.3.3 Direct invocation of SQL

Direct invocation of SQL is a method of executing SQL-statements directly. In direct invocation of SQL, the following are implementation-defined:

— The method of invoking SQL-statements.

— The method of raising conditions that result from the execution of such statements.

— The method of accessing the diagnostics information that results from the execution of such statements.

— The method of returning the results.

5.3.4 SQL-statements specified in [ISO9075-2]

The following are the classes of SQL-statements specified in [ISO9075-2]:

— SQL-schema statements, which can be used to create, alter, and drop schemas and the schema objects specified in [ISO9075-2].

— SQL-data statements, which can be used to perform queries, and insert, update and delete operations on tables. Some SQL-data statements contain the word “dynamic” in their names; they are not to be confused with SQL-dynamic statements.

— SQL-transaction statements, which can be used to set properties of, and initiate or terminate transactions.

— Two SQL-control statement (CALL and RETURN), which can be used to invoke a procedure and to specify a value to be returned by a function, respectively.

— SQL-connection statements, which can be used to initiate and terminate connections, and allow an SQL-client to switch from an SQL-session with one SQL-server to an SQL-session with another.

— SQL-session statements, which can be used to set some default values and other properties of an SQL-session.

— SQL-diagnostics statements, which get diagnostic information (from the diagnostics area).

— SQL-dynamic statements, which support the preparation and execution of dynamically generated SQL-statements, and obtaining information about them.

— SQL embedded exception declaration, which is converted to a statement in the host language.

For each SQL-statement that it defines, [ISO9075-2] specifies which SQL-statements will, if executed when no transaction is active, initiate a transaction and which will not.

For each SQL-statement, [ISO9075-2] specifies whether:

— It may be embedded in a host language.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

— It may be dynamically prepared and executed. Any preparable SQL-statement can be executed immediately, with the exception of those that fetch data into a descriptor area.

— It may be executed directly.

5.4 ISO/IEC 9075-3: Call Level Interface (SQL/CLI)

[ISO9075-3] specifies a method of binding between a program, written in one of a number of programming languages, and an SQL-implementation. The effect is functionally equivalent to dynamic SQL, specified in [ISO9075-2].

Procedures (routines) are specified that can be used to:

— Allocate and free resources (descriptor, or communication areas).

— Initiate, control, and terminate SQL-connections between SQL-client and SQL-servers.

— Cause the execution of SQL-statements, including the preparation of statements for subsequent execution.

— Obtain diagnostic information.

— Obtain information about the SQL-implementation, for example, the SQL-servers to which the SQL-client may be able to connect.

An important difference between CLI and Bindings is that, in the context of the latter, there is only one SQL-environment, whereas, in the context of CLI, a number of SQL-environments can be initiated and managed independently. Consequently, although an SQL-environment is defined to be simply a set of circumstances with various features, the term is used in CLI to refer to the current state (descriptor) of one SQL-environment, possibly among many. Thus, the term is used to mean the session between an application (SQL-agent) and an SQL-client (not to be confused with the SQL-session — referred to in CLI as the SQL-connection — between SQL-client and SQL-server).

5.5 ISO/IEC 9075-4: Persistent Stored Modules (SQL/PSM)

[ISO9075-4] makes SQL computationally complete by specifying the syntax and semantics of additional SQL-statements.

Those include facilities for:

— The specification of statements to direct the flow of control.

— The assignment of the result of expressions to variables and parameters.

— The specification of condition handlers that allow compound statements to deal with various conditions that may arise during their execution.

— The specification of statements to signal and resignal conditions.

— The declaration of standing SQL-server cursors.

— The declaration of local variables.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

It also defines Information Schema tables that contain schema information describing SQL-server modules.

5.5.1 SQL-statements specified in [ISO9075-4]

The following are the broad classes of SQL-statements specified in [ISO9075-4]:

— Additional SQL-control statements, which may be used to control the execution of an SQL routine.

— SQL-control declarations, which may be used to declare variables and exception handlers.

— Additional SQL-diagnostics statements, which may be used to signal exceptions.

— Additional SQL-schema statements, which may be used to create and drop modules.

5.6 ISO/IEC 9075-9: Management of External Data (SQL/MED)

[ISO9075-9] defines facilities that allow Database Language SQL to support management of external data through the use of foreign tables and datalink data types.

These include facilities for defining:

— Foreign servers.

— Foreign-data wrappers.

— Foreign tables.

and routines that can be used by an SQL-server to communicate with a foreign-data wrapper to interact with an external data source.

5.7 ISO/IEC 9075-10: Object Language Bindings (SQL/OLB)

[ISO9075-10] defines facilities for the embedding of SQL statements in Java programs.

5.8 ISO/IEC 9075-11: Information and Definition Schemas (SQL/Schemata)

[ISO9075-11], Information and Definition Schemas, specifies two schemas, the Information Schema (INFORMATION_SCHEMA) and the Definition Schema (DEFINITION_SCHEMA). The views of the Information Schema enable an application to learn the names of persistent database objects, such as tables, views, columns, and so forth. These views are defined in terms of the base tables of the Definition Schema. The only purpose of the Definition Schema is to provide a data model to support the Information Schema and to assist understanding. An SQL-implementation need do no more than simulate the existence of the Definition Schema, as viewed through the Information Schema views.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

5.9 ISO/IEC 9075-13: Java Routines & Types Using the Java Programming Language (SQL/JRT)

[ISO9075-13] defines facilities that allow Database Language SQL to enable invocations of static methods written in the Java programming language as SQL-invoked routines, and to use classes defined in the Java programming language as SQL structured types.

These include:

— Extensions to the definition, manipulation, and invocation of SQL-invoked routines.

— Extensions to the definition and manipulation of user-defined types.

— New built-in procedures.

5.10 ISO/IEC 9075-14: XML-Related Specifications (SQL/XML)

[ISO9075-14] defines facilities that allow Database Language SQL to enable creation and manipulation of XML documents.

These include:

— A new predefined type, XML.

— New built-in operators to create and manipulate values of the XML type.

— Rules for mapping tables, schemas, and catalogs to XML documents.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6 Notation and conventions used in other parts of ISO/IEC 9075

The notation and conventions defined in this clause are used in the other parts of ISO/IEC 9075 except where more appropriate ones are defined locally.

6.1 Notation taken from [ISO10646]

The notation for the representation of UCS code points is defined in [UCS], Subclause 6.5, "Short identifiers for characters".

In this International Standard, this notation is used only to unambiguously identify characters and is not meant to imply a specific encoding for any implementation's use of that character.

6.2 Notation provided in this International Standard

The syntactic notation used in ISO/IEC 9075 is an extended version of BNF ("Backus Normal Form" or "Backus Naur Form").

In a BNF language definition, each syntactic element, known as a *BNF nonterminal symbol*, of the language is defined by means of a *production rule*. This defines the element in terms of a formula consisting of the characters, character strings, and syntactic elements that can be used to form an instance of it.

In the version of BNF used in ISO/IEC 9075, the following symbols have the meanings shown:

Table 2 — Symbols used in BNF

Symbol	*Meaning*
`< >`	A character string enclosed in angle brackets is the name of a syntactic element (BNF nonterminal) of the SQL language.
`::=`	The definition operator is used in a production rule to separate the element defined by the rule from its definition. The element being defined appears to the left of the operator and the formula that defines the element appears to the right.
`[ ]`	Square brackets indicate optional elements in a formula. The portion of the formula within the brackets may explicitly specified or may be omitted.
`{ }`	Braces group elements in a formula. The portion of the formula within the braces shall be explicitly specified.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Symbol	*Meaning*
|	The alternative operator. The vertical bar indicates that the portion of the formula following the bar is an alternative to the portion preceding the bar. If the vertical bar appears at a position where it is not enclosed in braces or square brackets, it specifies a complete alternative for the element defined by the production rule. If the vertical bar appears in a portion of a formula enclosed in braces or square brackets, it specifies alternatives for the contents of the innermost pair of such braces or brackets.
. . .	The ellipsis indicates that the element to which it applies in a formula may be repeated any number of times. If the ellipsis appears immediately after a closing brace "}", then it applies to the portion of the formula enclosed between that closing brace and the corresponding opening brace "{". If an ellipsis appears after any other element, then it applies only to that element. In Syntax Rules, Access Rules, General Rules, and Conformance Rules, a reference to the *n*-th element in such a list assumes the order in which these are specified, unless otherwise stated.
! !	Introduces normal English text. This is used when the definition of a syntactic element is not expressed in BNF.

Spaces are used to separate syntactic elements. Multiple spaces and line breaks are treated as a single space. Apart from those symbols to which special functions were given above, other characters and character strings in a formula stand for themselves. In addition, if the symbols to the right of the definition operator in a production consist entirely of BNF symbols, then those symbols stand for themselves and do not take on their special meaning.

Pairs of braces and square brackets may be nested to any depth, and the alternative operator may appear at any depth within such a nest.

A character string that forms an instance of any syntactic element may be generated from the BNF definition of that syntactic element by application of the following steps:

1) Select any one option from those defined in the right hand side of a production rule for the element, and replace the element with this option.

2) Replace each ellipsis and the object to which it applies with one or more instances of that object.

3) For every portion of the string enclosed in square brackets, either delete the brackets and their contents or change the brackets to braces.

4) For every portion of the string enclosed in braces, apply step 1) through step 5) to the substring between the braces, then remove the braces.

5) Apply steps step 1) through step 5) to any BNF non-terminal symbol that remains in the string.

The expansion or production is complete when no further non-terminal symbols remain in the character string.

The left normal form derivation of a character string *CS* in the source language character set from a BNF non-terminal *NT* is obtained by applying step 1) through step 5) above to *NT*, always selecting in step 5) the leftmost BNF nonterminal.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6.3 Conventions

6.3.1 Specification of syntactic elements

Syntactic elements are specified in terms of:

— **Function**: A short statement of the purpose of the element.

— **Format**: A BNF definition of the syntax of the element.

— **Syntax Rules**: A specification in English of the syntactic properties of the element, or of additional syntactic constraints, not expressed in BNF, that the element shall satisfy, or both.

— **Access Rules**: A specification in English of rules governing the accessibility of schema objects that shall hold before the General Rules may be successfully applied.

— **General Rules**: A specification in English of the run-time effect of the element. Where more than one General Rule is used to specify the effect of an element, the required effect is that which would be obtained by beginning with the first General Rule and applying the Rules in numeric sequence unless a Rule is applied that specifies or implies a change in sequence or termination of the application of the Rules. Unless otherwise specified or implied by a specific Rule that is applied, application of General Rules terminates when the last in the sequence has been applied.

— **Conformance Rules**: A specification of how the element shall be supported for conformance to SQL.

The scope of notational symbols is the Subclause in which those symbols are defined. Within a Subclause, the symbols defined in Syntax Rules, Access Rules, or General Rules can be referenced in other rules provided that they are defined before being referenced.

6.3.2 Specification of the Information and Definition Schemata

The objects of the Information and Definition Schemata in ISO/IEC 9075 are specified in terms of:

— **Function**: A short statement of the purpose of the definition.

— **Definition**: A definition, in SQL, of the object being defined.

— **Description**: A specification of the run-time value of the object, to the extent that this is not clear from the definition.

— **Conformance Rules**: A specification of how the element shall be supported for conformance to SQL.

The only purpose of the view definitions in the Information Schema is to specify the contents of those viewed tables. The actual objects on which these views are based are implementation-dependent.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6.3.3 Use of terms

6.3.3.1 Syntactic containment

Let <A>, <B>, and <C> be syntactic elements; let *A1*, *B1*, and *C1* respectively be instances of <A>, <B>, and <C>.

In a Format, <A> is said to *immediately contain* <B> if <B> appears on the right-hand side of the BNF production rule for <A>. An <A> is said to *contain* or *specify* <C> if <A> immediately contains <C> or if <A> immediately contains a <B> that contains <C>.

In SQL language, *A1* is said to *immediately contain B1* if <A> immediately contains <B> and *B1* is part of the text of *A1*. *A1* is said to *contain* or *specify C1* if *A1* immediately contains *C1* or if *A1* immediately contains *B1* and *B1* contains *C1*. If *A1* contains *C1*, then *C1* is *contained in A1* and *C1* is *specified by A1*.

A1 is said to contain *B1 with an intervening* <C> if *A1* contains *B1* and *A1* contains an instance of <C> that contains *B1*. *A1* is said to contain *B1 without an intervening* <C> if *A1* contains *B1* and *A1* does not contain an instance of <C> that contains *B1*.

A1 simply contains B1 if *A1* contains *B1* without an intervening instance of <A> or an intervening instance of <B>.

A1 directly contains B1 if *A1* contains *B1* without an intervening <subquery>, <multiset value constructor by query>, <table value constructor by query>, <array value constructor by query>, <within group specification>, or <set function specification> that is not an <ordered set function>.

If <A> contains <B>, then <B> is said to be *contained in* <A> and <A> is said to be a *containing* production symbol for <B>. If <A> simply contains <B>, then <B> is said to be *simply contained in* <A> and <A> is said to be a *simply containing* production symbol for <B>.

A1 is the *innermost* <A> satisfying a condition *C* if *A1* satisfies *C* and *A1* does not contain an instance of <A> that satisfies *C*. *A1* is the *outermost* <A> satisfying a condition *C* if *A1* satisfies *C* and *A1* is not contained in an instance of <A> that satisfies *C*.

If <A> contains a <table name> that identifies a view that is defined by a <view definition> *V*, then <A> is said to *generally contain* the <query expression> contained in *V*. If <A> contains a <query name> that identifies a <query expression> *QE*, then <A> is said to *generally contain QE*. If <A> contains a <routine invocation> *RI*, then <A> is said to *generally contain* the routine bodies of all SQL-invoked routines in the set of subject routines of *RI*. If <A> contains <B>, then <A> generally contains <B>. If <A> generally contains <B> and <B> generally contains <C>, then <A> generally contains <C>.

NOTE 5 — The “set of subject routines of a <routine invocation>” is defined in Subclause 10.4, “<routine invocation>”, in [ISO9075-2].

In a Format, the verb “to be” (including all its grammatical variants, such as “is”) is defined as follows: <A> is said to be <B> if there exists a BNF production rule of the form `<A> ::= <B>`. If <A> is <B> and <B> is <C>, then <A> is <C>. If <A> is <C>, then <C> is said to *constitute* <A>. In SQL language, *A1* is said to be *B1* if <A> is <B> and the text of *A1* is the text of *B1*. Conversely, *B1* is said to *constitute A1* if *A1* is *B1*.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6.3.3.2 Terms denoting rule requirements

In the Syntax Rules, the term *shall* defines conditions that are required to be true of syntactically conforming SQL language. When such conditions depend on the contents of one or more schemas, they are required to be true just before the actions specified by the General Rules are performed. The treatment of language that does not conform to the SQL Formats and Syntax Rules is implementation-dependent. If any condition required by Syntax Rules is not satisfied when the evaluation of Access or General Rules is attempted and the implementation is neither processing non-conforming SQL language nor processing conforming SQL language in a non-conforming manner, then an exception condition is raised: *syntax error or access rule violation*.

In the Access Rules, the term *shall* defines conditions that are required to be satisfied for the successful application of the General Rules. If any such condition is not satisfied when the General Rules are applied, then an exception condition is raised: *syntax error or access rule violation*.

In the Conformance Rules, the term *shall* defines conditions that are required to be true if the named Feature or Features are not required to be supported.

6.3.3.3 Rule evaluation order

A conforming implementation is not required to perform the exact sequence of actions defined in the General Rules, provided its effect on SQL-data and schemas, on host parameters and host variable, and on SQL parameters and SQL variables is identical to the effect of that sequence. The term *effectively* is used to emphasize actions whose effect might achieved in other ways by an implementation.

The Syntax Rules and Access Rules for contained syntactic elements are effectively applied at the same time as the Syntax Rules and Access Rules for the containing syntactic elements. The General Rules for contained syntactic elements are effectively applied before the General Rules for the containing syntactic elements.

Where the precedence of operators is determined by the Formats of ISO/IEC 9075 or by parentheses, those operators are effectively applied in the order specified by that precedence.

Where the precedence is not determined by the Formats or by parentheses, effective evaluation of expressions is *generally* performed from left to right. However, it is implementation-dependent whether expressions are *actually* evaluated left to right, particularly when operands or operators might cause conditions to be raised or if the results of the expressions can be determined without completely evaluating all parts of the expression.

In general, if some syntactic element contains more than one other syntactic element, then the General Rules for contained elements that appear earlier in the production for the containing syntactic element are applied before the General Rules for contained elements that appear later.

For example, in the production:

```
<A> ::=
  <B> <C>
```

the Syntax Rules and Access Rules for <A>, <B>, and <C> are effectively applied simultaneously. The General Rules for <B> are applied before the General Rules for <C>, and the General Rules for <A> are applied after the General Rules for both <B> and <C>.

> NOTE 6 — There are exceptions to this practice. For example, in [ISO9075-2], Subclause 13.5, "<SQL procedure statement>", illustrates one kind of exception to this general rule for General Rules. In that Subclause, the General Rules of the particular contained statement (e.g., an <insert statement>) are invoked only when a General Rule in Subclause 13.5, "<SQL procedure statement>", explicitly states that the General Rules of the contained statement are to be evaluated.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

If the result of an expression or search condition is not dependent on the result of some part of that expression or search condition, then that part of the expression or search condition is said to be *inessential*. An invocation of an SQL-invoked function that is deterministic and that does not possibly modify SQL-data is considered to be *inessential*. It is implementation-defined whether an invocation of an SQL-invoked function that is nondeterministic or possibly modifies SQL-data is considered to be *essential* or *inessential*.

If an Access Rule pertaining to an inessential part is not satisfied, then the *syntax error or access rule violation* exception condition is raised regardless of whether or not the inessential parts are actually evaluated. If evaluation of an inessential part would cause an exception condition to be raised, then it is implementation-dependent whether or not that exception condition is raised.

During the computation of the result of an expression, the SQL-implementation may produce one or more *intermediate results* that are used in determining that result. The declared type of a site that contains an intermediate result is implementation-dependent.

6.3.3.4 Conditional rules

A conditional rule is specified with “If” or “Case” conventions. A rule specified with “Case” conventions includes a list of conditional subrules using “If” conventions. The first such “If” subrule whose condition is true is the effective subrule of the “Case” rule. The last subrule of a “Case” rule may specify “Otherwise”, in which case it is the effective subrule of the “Case” rule if no preceding “If” subrule in the “Case” rule is satisfied. If the last subrule does not specify “Otherwise”, and if there is no subrule whose condition is true, then there is no effective subrule of the “Case” rule.

6.3.3.5 Syntactic substitution

In the Syntax and General Rules, the phrase “ *X* is implicit” indicates that the Syntax and General Rules are to be interpreted as if the element *X* had actually been specified. Within the Syntax Rules of a given Subclause, it is known whether the element was explicitly specified or is implicit.

In the Syntax and General Rules, the phrase “the following <x> is implicit: *Y*” indicates that the Syntax and General Rules are to be interpreted as if a syntactic element <x> containing *Y* had actually been specified.

In the Syntax Rules and General Rules, the phrase “*former* is equivalent to *latter*” indicates that the Syntax Rules and General Rules are to be interpreted as if all instances of *former* in the element had been instances of *latter*.

If a BNF nonterminal is referenced in a Subclause without specifying how it is contained in a BNF production that the Subclause defines, then

Case:

— If the BNF nonterminal is itself defined in the Subclause, then the reference shall be assumed to be to the occurrence of that BNF nonterminal on the left side of the defining production.

— Otherwise, the reference shall be assumed to be to a BNF production in which the particular BNF nonterminal is immediately contained.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6.3.3.6 Other terms

Some Syntax Rules define terms, such as *T1*, to denote named or unnamed tables. Such terms are used as table names or correlation names. Where such a term is used as a correlation name, it does not imply that any new correlation name is actually defined for the denoted table, nor does it affect the scopes of any actual correlation names.

An SQL-statement *S1* is said to be executed as a *direct result* of the execution an SQL-statement *S2* if *S2* is a <call statement> *CS* and *S1* is the outermost SQL-statement contained in the SQL-invoked routine that is the subject routine of the <routine invocation> contained in *CS*.

An SQL-statement *S1* is said to be executed as an *indirect result* of executing an SQL-statement *S2* if *S1* is executed in a trigger execution context that has been activated by the execution of *S2*.

A value *P is part of* a value *W* if and only if:

— *W* is a table and *P* is a row of *W*.

— *W* is a row and *P* is a field of *W*.

— *W* is a collection and *P* is an element of *W*.

— *P* is a part of some value that is a part of *W*.

If a value has parts, then it follows that an instance of that value has parts; hence the site it occupies has parts, each of which is also a site.

An item *X is part of* an item *Y* if and only if:

— *Y* is a row and *X* is a column of *Y*.

— *Y* is a <routine invocation> and *X* is an SQL parameter of *Y*.

— *Y* is a user-defined type instance and *X* is an attribute of *Y*.

— There exists an item *X2* such that *X* is a part of *X2* and *X2* is a part of *Y*.

Another part of ISO/IEC 9075 may define additional terms that are used in that part only.

6.3.3.7 Exceptions

Except where otherwise specified, the phrase “an exception condition is raised:”, followed by the name of a condition, is used in General Rules and elsewhere to indicate that:

— The execution of a statement is unsuccessful.

— The application of the General Rules of Subclauses not specified in Subclause 6.3.3.8, “General Rules not terminated on exception conditions”, may be terminated.

— Diagnostic information is to be made available.

— Execution of the statement is to have no effect on SQL-data or schemas.

The effect on any assignment target and SQL descriptor area of an SQL-statement that terminates with an exception condition, unless explicitly defined by ISO/IEC 9075, is implementation-dependent.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

The phrase “a completion condition is raised”, followed by a colon and the name of a condition, is used in General Rules and elsewhere to indicate that application of General Rules is not terminated and diagnostic information is to be made available; unless an exception condition is also raised, the execution of the SQL-statement is successful.

If more than one condition could have occurred as a result of a statement, it is implementation-dependent whether diagnostic information pertaining to more than one condition is made available. See Subclause 4.29.2, “Status parameters”, in [ISO9075-2], for rules regarding precedence of status parameter values.

6.3.3.8 General Rules not terminated on exception conditions

The General Rules of the following Subclauses save and restore context and so, except where explicitly noted, are not terminated where an exception condition is raised:

— Subclause 10.4, “<routine invocation>”, in [ISO9075-2].

— Subclause 13.5, “<SQL procedure statement>”, in [ISO9075-2].

— Subclause 15.19, “Execution of triggers”, in [ISO9075-2].

— Subclause 22.1, “<direct SQL statement>”, in [ISO9075-2].

— Subclause 14.1, “<compound statement>”, in [ISO9075-4].

— Subclause 14.2, “<handler declaration>”, in [ISO9075-4].

6.3.4 Descriptors

A descriptor is a conceptual structured collection of data that defines an object of a specified type. The concept of descriptor is used in specifying the semantics of SQL. It is not necessary that any descriptor exist in any particular form in any SQL-environment.

Some SQL objects cannot exist except in the context of other SQL objects. For example, columns cannot exist except within the context of tables. Each such object is independently described by its own descriptor, and the descriptor of an enabling object (*e.g.*, table) is said to *include* the descriptor of each enabled object (*e.g.*, column or table constraint). Conversely, the descriptor of an enabled object is said to *be included in* the descriptor of an enabling object. The descriptor of some object *A* *generally includes* the descriptor of some object *C* if the descriptor of *A* includes the descriptor of some object *B* and the descriptor of *B* generally includes the descriptor of *C*.

In other cases, certain SQL objects cannot exist unless some other SQL object exists, even though there is no inclusion relationship. For example, SQL does not permit an assertion to exist if some table referenced by the assertion does not exist. Therefore, an assertion descriptor *is dependent on* or *depends on* one or more table descriptors (equivalently, an assertion *is dependent on* or *depends on* one or more tables). In general, a descriptor *D1* can be said to depend on, or to be dependent on, some descriptor *D2*.

The descriptor of some object *A* *generally depends on* or *is generally dependent on* the descriptor of some object *C* if the descriptor of *A* depends on the descriptor of some object *B* and the descriptor of *B* generally depends on the descriptor of *C*.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

The execution of an SQL-statement may result in the creation of many descriptors. An SQL object that is created as a result of an SQL-statement may depend on other descriptors that are only created as a result of the execution of that SQL statement.

NOTE 7 — This is particularly relevant in the case of the <schema definition> SQL-statement. A <schema definition> may, for example, contain many <table definition>s that in turn contain <table constraint>s. A single <table constraint> in one <table definition> may reference a second table being created by a separate <table definition> which itself may contain a reference to the first table. The dependencies of each table on the descriptors of the other are valid provided that all necessary descriptors are created during the execution of the complete <schema definition>.

There are two ways of indicating dependency of one SQL object on another. In many cases, the descriptor of the dependent SQL object is said to "include the name of" the SQL object on which it is dependent. In this case "the name of" is to be understood as meaning "sufficient information to identify the descriptor of". Alternatively, the descriptor of the dependent SQL object may be said to include text (*e.g.*, <query expression>, <search condition>) of the SQL object on which it is dependent. However, in such cases, whether the implementation includes actual text (with defaults and implications made explicit) or its own style of parse tree is irrelevant; the validity of the descriptor is clearly "dependent on" the existence of descriptors of objects that are referenced in it.

The statement that a column "is based on" a domain, is equivalent to a statement that a column "is dependent on" that domain.

An attempt to destroy an SQL object, and hence its descriptor, may fail if other descriptors are dependent on it, depending on how the destruction is specified. Such an attempt may also fail if the descriptor to destroyed is included in some other descriptor. Destruction of a descriptor results in the destruction of all descriptors included in it, but has no effect on descriptors on which it is dependent.

The implementation of some SQL objects described by descriptors requires the existence of objects not specified by this International Standard. Where such objects are required, they are effectively created whenever the associated descriptor is created and effectively destroyed whenever the associated descriptor is destroyed.

6.3.5 Relationships of parts within ISO/IEC 9075

Parts of ISO/IEC 9075 other than this part of ISO/IEC 9075 and [ISO9075-2] depend on [ISO9075-1] and [ISO9075-2] and its Technical Corrigenda and are referenced as *incremental parts*. Each incremental part is to be used as though it were merged with the text of ISO/IEC 9075. This Subclause describes the conventions used to specify the merger.

The merger described also accounts for the Technical Corrigenda that have been published to correct ISO/IEC 9075. This accommodation is typically indicated by the presence of a phrase like "in the Technical Corrigenda" or "in the TC".

6.3.5.1 New and modified Clauses, Subclauses, and Annexes

Where a Clause (other than Clause 1, "Scope", and Clause 2, "Normative references"), Subclause, or Annex in any incremental part of ISO/IEC 9075 has a name identical to a Clause, Subclause, or Annex in [ISO9075-2], [ISO9075-3], [ISO9075-4], or [ISO9075-11], unless the incremental part is itself [ISO9075-3], [ISO9075-4], or [ISO9075-11], it supplements the Clause, Subclause, or Annex, respectively, in [ISO9075-2], [ISO9075-3], [ISO9075-4], or [ISO9075-11], regardless of whether or not the number or letter of the Clause, Subclause,

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

or Annex corresponds. It typically does so by adding or replacing paragraphs, Format items, Tables, Figures, Rules, or Notes.

In each incremental part, the relationships between each Clause, Subclause, and Annex in that incremental part and the corresponding Clause, Subclause, or Annex in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] are shown by:

— In the incremental part: A statement of the form "*This Object modifies Object nn.nn, "xxxxx", in ISO/IEC 9075-n.*" immediately follows the Object title.

— In the part referenced by the statement inserted in the preceding step: A statement of the form "*This Object is modified by Object mm.mmm, "yyyyy", in ISO/IEC 9075-m.*" immediately follows the Object title.

— The Object can be a Clause, a Subclause, or an Annex.

Where a Clause, Subclause, or Annex in an incremental part has a name that is not identical to the name of some Clause, Subclause, or Annex in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], it provides language specification particular to that incremental part. A Subclause that is part of a Clause or Subclause identified as new is inherently new and is not marked.

The Clauses, Subclauses, and Annexes in each incremental part appear in the order in which they are intended to appear in the merged document. In the absence of other explicit instructions regarding its placement, any new Clause, Subclause, or Annex is to be positioned as follows: Locate the prior Clause, Subclause, or Annex in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] whose name is identical to the name of a corresponding Clause, Subclause, or Annex that appears in the incremental part of ISO/IEC 9075. The new Clause, Subclause, or Annex shall immediately follow that Clause, Subclause, or Annex. If there are multiple new Clauses, Subclauses, or Annexes with no intervening Clause, Subclause, or Annex that modifies an existing Clause, Subclause, or Annex, then those new Clauses, Subclauses, or Annexes appear in order, following the prior Clause, Subclause, or Annex whose name was matched.

When an incremental part performs a modification to a Clause, Subclause, or Annex in [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], then the modifications are applied in the following sequence:

1) All modifications to [ISO9075-3] from the incremental part.

2) All modifications to [ISO9075-4] from the incremental part.

3) All modifications to [ISO9075-11] from the incremental part.

4) All modifications to [ISO9075-2] from [ISO9075-11], (including all modifications that were added, augmented, or replaced as a result of step 3)).

5) All modifications to [ISO9075-2] from [ISO9075-3], (including all modifications that were added, augmented, or replaced as a result of step 1)).

6) All modifications to [ISO9075-2] from [ISO9075-4], (including all modifications that were added, augmented, or replaced as a result of step 2)).

7) All modifications to [ISO9075-2] from the incremental part. Note that modifications in this final step may augment or replace modifications applied as a result of step 4), step 5), and step 6).

Modifications to one or more of [ISO9075-2], [ISO9075-3], [ISO9075-4], and [ISO9075-11] by more than one incremental part do not interact. The modifications made by an incremental part only have influence on the language specification of that part and those specifications are not influenced by modifications made by any other incremental part.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

6.3.5.2 New and modified tables and figures

Where a Table in any incremental part of ISO/IEC 9075 has a name identical to a Table in [ISO9075-2], [ISO9075-3], [ISO9075-4], or [ISO9075-11], unless the incremental part is itself [ISO9075-3], [ISO9075-4], or [ISO9075-11], it supplements the Table in [ISO9075-2], [ISO9075-3], [ISO9075-4], or [ISO9075-11], regardless of whether or not the number of the Table corresponds.

The rows in modified tables are generally new rows to be effectively inserted into the corresponding table, though in rare cases a row already in a table is effectively replaced by a row in the table in the incremental part. Such replacement is required wherever the value in the first column of the corresponding table is the same.

Where a Figure in any incremental part of ISO/IEC 9075 has a name identical to a Figure in [ISO9075-2], [ISO9075-3], [ISO9075-4], or [ISO9075-11], unless the incremental part is itself [ISO9075-3], [ISO9075-4], or [ISO9075-11], it supplements the Figure in [ISO9075-2], [ISO9075-3], [ISO9075-4], or [ISO9075-11], regardless of whether or not the number of the Figure corresponds.

When a Table or Figure in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is modified or replaced by a Table or Figure in some incremental Part, in that incremental part, a statement of the form "*This Object modifies Object nn.nn, "xxxxx", in ISO/IEC 9075-n.*" is typically found in a Rule or paragraph that appears a little above the modifying Object.

When a Table or Figure in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is modified or replaced by a Table or Figure in some incremental Part, the caption of the Table or Figure in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is preceded by a small red box containing the number of the incremental Part that contains the modifying or replacing Table or Figure.

Where a Table or Figure in an incremental part has a name that is not identical to the name of some Table or Figure in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], it provides specifications particular to that incremental part.

6.3.5.3 Functions

In a modified Subclause, the Function from the modifying Subclause completely replaces the Function in the original Subclause.

6.3.5.4 New and modified Format items

In modified Clauses and Subclauses, a Format item that defines a BNF nonterminal symbol (that is, the BNF nonterminal symbol that appears on the left-hand side of the `::=` mark) either modifies a Format item whose definition appears in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], or replaces a Format item whose definition appears in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], or defines a new Format item that does not have a definition at all in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11].

Those Format items in the incremental part that modify a Format item whose definition appears in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] are identified by the existence of a "Format comment" such as:

```
<modified item> ::=
  !! All alternatives from ISO/IEC 9075-2
| <new alternative>
```


Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS Not for Resale

By contrast, Format items that completely replace Format items in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] have BNF nonterminal symbols identical to BNF nonterminal symbols of Format items in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], but do not state that they include any alternatives from [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11].

New Format items that have no correspondence to any Format item in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] are not distinguished in the incremental part.

Format items in new Subclauses are unmarked.

When a Format item in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is modified or replaced by a Format item in some incremental Part, the Format item in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is indicated by the presence of a small red box, immediately preceding the name of the BNF nonterminal being defined by the Format item, containing the number of the incremental Part that contains the modifying or replacing Format item.

6.3.5.5 New and modified paragraphs and rules

In modified Clauses and Subclauses, each paragraph or Rule is marked to indicate whether it is a modification of a paragraph or Rule in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11], or is a new paragraph or Rule added by this incremental part.

Modifications of paragraphs or Rules are identified by the inclusion of an indicative phrase enclosed in a box.

Replace the 5th paragraph means that the following text is to replace the fifth paragraph of the corresponding Subclause in the part identified in the statement following the Subclause title.

Replace SR6)b)ii) means that the following text is to replace Syntax Rule 6)b)ii) of the corresponding Subclause in the part identified in the statement following the Subclause title.

Augments SR3) means that the following text is to extend or enhance Syntax Rule 3). In most instances, the augmentation is the addition of a new alternative meant to support new syntax.

"Replace" without qualification means the entire paragraph or rule at the level that is indicated is to be replaced. For example, Replace SR6)b) means that Syntax Rule 6)b) and any subrules that it may contain is to be replaced. If the term *lead text of* is included, then the replacement is for the text of the paragraph or rule up to, but not including, the first bulleted, dashed, or enumerated list item that follows. For example, Replace the lead text of SR1)b) means that, in the following rule, only the italicized text is to be replaced.

1) Lead text for Syntax Rule 1)

   a) Text for Syntax Rule 1)a)

   b) *Lead text for Syntax Rule 1)b).*

      Case:

      i) Text for Syntax Rule 1)b)i)

      ii) Text for Syntax Rule 1)b)ii)

   c) Text for Syntax Rule 1)c)

--`,,```,,,,````-`-`,,`,,`,`,,`---
© ISO/IEC 2008 – All rights reserved
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

New paragraphs or Rules in an incremental part are marked to indicate where they are to be inserted.

Insert before 2nd paragraph means that the following text is to read as though it were inserted immediately before the second paragraph of the corresponding Subclause in the part identified in the statement following the Subclause title.

Insert before GR4) means that the following text is to read as though it were inserted immediately before General Rule 4) of the corresponding Subclause in the part identified in the statement following the Subclause title.

If no specific insertion point is indicated, as in Insert this paragraph or Insert this GR, then the following text is to be read as though it were appended at the end of the appropriate section (the General Rules, for example) of the corresponding Subclause in the part identified in the statement following the Subclause title.

Paragraphs in each Subclause are numbered from 1 (one), and bulleted, dashed, or enumerated lists are counted as part of the same paragraph as the text of the paragraph. Numbered notes are not counted as paragraphs.

In such indications, “SR” is used to mean “Syntax Rule”, “AR” is used to mean “Access Rule”, “GR” is used to mean “General Rule”, and “CR” is used to mean “Conformance Rule”. “Desc.” is used to mean “Description” and “Func.” is used to mean “Function”.

When a Rule or paragraph in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is modified or replaced by a Rule or paragraph in some incremental Part, the Rule or paragraph in [ISO9075-2] and/or [ISO9075-3] and/or [ISO9075-4] and/or [ISO9075-11] is indicated by the presence of a small red box, immediately preceding the text of the Rule or paragraph, containing the number of the incremental Part that contains the modifying or replacing Rule or paragraph.

All paragraphs, Format items, and Rules in new Clauses or Subclauses are also new and are therefore unmarked.

6.3.5.6 Modified annexes

In a modified Annex, the lead text of the paragraphs of the modifying Annex are discarded.

NOTE 8 — This results in only the tables and the bulleted or enumerated list items being retained.

All other text and/or tables are merged into the text and/or tables of the original Annex, taking account of any grouping and sorting implied by the initial paragraphs of the original Annex.

NOTE 9 — References to Clauses and Subclauses in the text will be adjusted to the values that the referenced Clauses and Subclauses have acquired in the merged text.

6.3.6 Subclauses used as subroutines

In various parts of this International Standard, some Subclauses are defined without explicit syntax to invoke their semantics. Such Subclauses, called subroutine Subclauses, typically factor out rules that are required by one or more other Subclauses and are intended to be invoked by the rules of those other Subclauses. In a few cases, these Subclauses not having explicit syntax are intended to be invoked by other standards and/or through the use of implementation-defined mechanisms.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

In other words, the rules of these Subclauses behave as though they were a sort of definitional "subroutine" that is invoked by other Subclauses, other standards, or implementation-defined mechanisms. These subroutine Subclauses are typically specified in a manner that requires information to be passed to them from their invokers. The information that must be passed is represented as parameters of these subroutine Subclauses, and that information must be passed in the form of arguments provided by the invokers of these subroutine Subclauses.

Every invocation of a subroutine Subclause must explicitly provide information for every required parameter of the subroutine Subclause being invoked. If a subroutine Subclause is invoked without the provision of information for every required parameter, the effects are implementation-dependent.

6.3.7 Index typography

In the Indexes to the parts of ISO/IEC 9075, the following conventions are used:

— An index entry in **boldface** indicates the page where the word, phrase, or BNF nonterminal is defined.

— An index entry in *italics* indicates a page where the BNF nonterminal is used in a Format.

— An index entry in neither boldface nor italics indicates a page where the word, phrase, or BNF nonterminal is not defined, but is used other than in a Format (for example, in a heading, Function, Syntax Rule, Access Rule, General Rule, Conformance Rule, Table, or other descriptive text).

6.3.8 Feature ID and Feature Name

Features are either *standard-defined features* or *implementation-defined features*.

Standard-defined features are defined in various parts of ISO/IEC 9075. Implementation-defined features are defined by SQL-implementations (see Subclause 8.4.8, “Extensions and options”).

Features are referenced by a Feature ID and by a Feature Name. A Feature ID value comprises either a letter and three digits or a letter, three digits, a hyphen, and one or two additional digits. Feature ID values containing a hyphen and additional digits indicate “subfeatures” that help to define complete features, which are in turn indicated by Feature ID values without a hyphen.

Feature IDs whose letter is “V” are reserved for implementation-defined features.

The set of subfeatures of a feature is intended to itemize noteworthy special cases of the feature, without necessarily being exhaustive of the feature. In addition, subfeatures of a feature are not necessarily mutually exclusive; in some cases, one subfeature may subsume another subfeature.

Only entire features are used to specify conformance requirements.

Standard-defined features are either mandatory features of a part, or they are optional features that are defined by Conformance Rules. The Feature ID of a standard-defined feature is stable and can be depended on to remain constant.

For convenience, all of the features defined in a part of ISO/IEC 9075 are collected together in a non-normative annex in the part in which they are defined. For example, the features defined in [ISO9075-2] can be found in Annex F, “SQL feature taxonomy”, in [ISO9075-2].

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Conformance Rules are generally placed in the Subclause that defines the BNF non-terminal that is controlled by the feature. In those circumstances where the use of a non-terminal is only controlled in a specific circumstance the Conformance Rule is placed where the non-terminal is used. This is also done in those circumstances where the fact that the use of the non-terminal is controlled by the feature can be deduced from an inspection of the Syntax Rules and the Conformance Rules of other Subclauses and the Conformance Rule is in principle redundant. As far as possible, other redundancy in the Conformance Rules is avoided.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

7 Annexes to the parts of ISO/IEC 9075

Every annex to every part of ISO/IEC 9075 is informative. The contents of each annex provides additional information, some of which restates that which is stated elsewhere in the normative text.

7.1 SQL conformance summary

[ISO9075-2] and every incremental part, other than [ISO9075-3], contains an Annex that lists every feature of ISO/IEC 9075, ordered by Feature number, including the number and name of every Subclause that contains a Conformance Rule associated with that feature.

7.2 Implementation-defined elements

Every part of ISO/IEC 9075 contains an Annex that lists every element of SQL and its processing that is specified in that part, and is permitted to differ between SQL-implementations, but is required to be specified by the implementor for each particular SQL-implementation.

7.3 Implementation-dependent elements

Every part of ISO/IEC 9075 contains an Annex that lists every element of SQL and its processing that is mentioned, but not specified in that part, and is thus permitted to differ between SQL-implementations, but is not required to be specified by the implementor for any particular SQL-implementation.

7.4 Deprecated features

[ISO9075-2] and every incremental part contains an Annex that lists every element of SQL and its processing that is specified in that part, but that may not be specified in some future revision of that part.

7.5 Incompatibilities with previous versions

[ISO9075-2] and every incremental part contains an Annex that lists every element of SQL and its processing that is specified in a previous version of that part, but that is not specified in the same way in the present version. The most frequent cause of such incompatibilities is the addition of reserved key words to the language, which invalidates their use in SQL language that conformed to an earlier version of ISO/IEC 9075.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

7.6 SQL feature taxonomy

[ISO9075-2] and every incremental part contains an Annex that lists every feature of ISO/IEC 9075 defined in that part, ordered by Feature number, specifying the name of that feature.

--`,,```,,,,````-`-`,,`,,`,`,,`---
© ISO/IEC 2008 – All rights reserved
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8 Conformance

Several different types of conformance may be claimed. Every claim of conformance shall include a minimum claim of conformance. In addition, conformance may be claimed for zero or more incremental parts, for zero or more optional features and for zero or more packages.

8.1 Minimum conformance

Every claim of conformance shall include a *claim of minimum conformance*, which is defined as a claim to meet the conformance requirements specified in [ISO9075-2] and [ISO9075-11]. A claim of minimum conformance shall include the statements required by the conformance requirements of [ISO9075-2] and [ISO9075-11].

SQL language that does not require more than a claim of minimum conformance is called *Core SQL*.

8.2 Conformance to parts

A claim of conformance to a part of ISO/IEC 9075 implies support of all mandatory features defined in that part.

In addition every claim of conformance to a part of ISO/IEC 9075 shall meet the conformance requirements specified in that part.

A claim of conformance to an incremental part of ISO/IEC 9075 shall include the claims required by the conformance requirements of that Part of ISO/IEC 9075 to which conformance is claimed.

8.3 Conformance to features

In addition to those features that are mandatory for conformance to any part of ISO/IEC 9075 (including [ISO9075-2] and [ISO9075-11]), any part of ISO/IEC 9075 may define optional features. These features are identified by Feature ID and controlled by Conformance Rules (see Subclause 6.3.8, “Feature ID and Feature Name”).

An optional feature *FEAT* is defined by relaxing selected Conformance Rules, as noted at the beginning of each Conformance Rule by the phrase “without Feature *FEAT*, “name of feature”, ...”. An application designates a set of SQL features that the application requires; the SQL language of the application shall observe the restrictions of all Conformance Rules except those explicitly relaxed for the required features. Conversely, conforming SQL-implementations shall identify which SQL features the SQL-implementation supports. An SQL-implementation shall process any application whose required features are a subset of the SQL-implementations supported features.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

A feature *FEAT1* may imply another feature *FEAT2*. An SQL-implementation that claims to support *FEAT1* shall also support each feature *FEAT2* implied by *FEAT1*. Conversely, an application need only designate that it requires *FEAT1*, and may assume that this includes each feature *FEAT2* implied by *FEAT1*. The list of features that are implied by other features is shown in a table named "Implied feature relationships of ..." in the Clause named "Conformance" in each part of ISO/IEC 9075. Note that some features imply multiple other features, and some features imply features defined in other parts. Every claim of conformance to an optional feature of ISO/IEC 9075 shall meet the conformance requirements of all parts to which conformance is claimed as if the Conformance Rules, which control the feature, did not exist.

The Syntax Rules and General Rules may define one SQL syntax in terms of another. Such transformations are presented to define the semantics of the transformed syntax, and are effectively performed after checking the applicable Conformance Rules. Transformations may use SQL syntax of one SQL feature to define another SQL feature. These transformations serve to define the behavior of the syntax, and do not have any implications for the feature syntax that is permitted or forbidden by the features so defined. A conforming SQL-implementation need only process the untransformed syntax defined by the Conformance Rules that are applicable for the set of features that the SQL-implementation claims to support, though with the semantics implied by the transformation.

8.4 Conformance to SQL packages

A package is a group of optional features of ISO/IEC 9075, which, taken together and in conjunction with Core SQL, provide a platform for some application area or implementation environment. A package is identified by a Package ID.

Every claim of conformance to a package of ISO/IEC 9075 shall meet the conformance requirements of all the parts and optional features which comprise the package.

This Subclause specifies several such packages. It is expected that packages may be specified elsewhere, outside of the scope of ISO/IEC 9075.

Table 3, "SQL Packages", contains a list of the packages of the SQL language defined by ISO/IEC 9075.

The "Package ID" column of Table 3, "SQL Packages", specifies the formal identification of the packages contained in the table. The Package ID is stable and can be depended on to remain constant.

The "Package Description" column of Table 3, "SQL Packages", contains a brief description of the package associated with the Package ID value.

Table 3 — SQL Packages

Package ID	Package Description
PKG001	Enhanced datetime facilities
PKG002	Enhanced integrity management
PKG004	PSM
PKG006	Basic object support

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Package ID	Package Description
PKG007	Enhanced object support
PKG008	Active database
PKG010	OLAP

8.4.1 Enhanced datetime facilities

The package called “Enhanced datetime facilities” comprises the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature F052, “Intervals and datetime arithmetic”

— Feature F411, “Time zone specification”

— Feature F555, “Enhanced seconds precision”

8.4.2 Enhanced integrity management

The package called “Enhanced integrity management” comprises the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature F191, “Referential delete actions”

— Feature F521, “Assertions”

— Feature F701, “Referential update actions”

— Feature F491, “Constraint management”

— Feature F671, “Subqueries in CHECK constraints”

— Feature T201, “Comparable data types for referential constraints”

— Feature T211, “Basic trigger capability”

— Feature T212, “Enhanced trigger capability”

— Feature T191, “Referential action RESTRICT”

8.4.3 PSM

The package called “PSM” comprises the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature T322, “Overloading of SQL-invoked functions and SQL-invoked procedures”

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

— Feature P001, “Stored modules”

— Feature P002, “Computational completeness”

— Feature P003, “Information Schema views”

8.4.4 Basic object support

The package called “Basic object support” comprises the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature S023, “Basic structured types”

— Feature S041, “Basic reference types”

— Feature S051, “Create table of type”>

— Feature S151, “Type predicate”

— Feature T041, “Basic LOB data type support”

8.4.5 Enhanced object support

The package called “Enhanced object support” comprises all of the features of the package called (Basic object support), plus the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature S024, “Enhanced structured types”

— Feature S043, “Enhanced reference types”

— Feature S071, “SQL-paths in function and type name resolution”

— Feature S081, “Subtables”

— Feature S111, “ONLY in query expressions”

— Feature S161, “Subtype treatment”

— Feature S211, “User-defined cast functions”

— Feature S231, “Structured type locators”

— Feature S241, “Transform functions”

8.4.6 Active database

The package called “Active database” comprises the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature T211, “Basic trigger capability”

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8.4.7 OLAP

The package called “OLAP” comprises the following features of the SQL language as specified in the SQL Feature Taxonomy Annex of the various parts of ISO/IEC 9075.

— Feature T431, “Extended grouping capabilities”

— Feature T611, “Elementary OLAP operators”

8.4.8 Extensions and options

An SQL-implementation may provide implementation-defined features that are additional to those specified by ISO/IEC 9075, and may add to the list of reserved words.

NOTE 10 — If additional words are reserved, it is possible that a conforming SQL-statement may not be processed correctly.

It is implementation-defined whether an SQL Flagger flags implementation-defined features.

NOTE 11 — An SQL Flagger may flag implementation-defined features using any Feature ID not defined by this International Standard. However, there is no guarantee that some future edition of this International Standard will not use such a Feature ID for a standard-defined feature.

NOTE 12 — The implementation-defined features flagged by an SQL Flagger may identify implementation-defined features from more than one SQL-implementation.

NOTE 13 — The allocation of a Feature ID to an implementation-defined feature may differ between SQL Flaggers.

NOTE 14 — An SQL-implementation may choose to, but is not required to, make an implementation-defined feature visible through the SQL_FEATURES view of the Information Schema.

An SQL-implementation may provide support for additional implementation-defined SQL-invoked routines or for implementation-defined argument values for SQL-invoked routines.

An SQL-implementation may provide user options to process non-conforming SQL statements or routine invocations. An SQL-implementation may provide user options to process SQL statements or routine invocations so as to produce a result different from that specified in the parts of ISO/IEC 9075.

It shall produce such results only when explicitly required by the user option. Additional extensions and options may be specified in incremental parts of ISO/IEC 9075.

8.5 SQL flagger

An SQL Flagger is an implementation-provided facility that is able to identify SQL language extensions, or other SQL processing alternatives, that may be provided by a conforming SQL-implementation (see Subclause 8.4.8, “Extensions and options”).

An SQL Flagger is intended to assist in the production of SQL language that is both portable and interoperable among different conforming SQL-implementations operating under different levels of this International Standard.

An SQL Flagger is intended to effect a static check of SQL language. There is no requirement to detect extensions that cannot be determined until the General Rules are evaluated.

An SQL-implementation need only flag SQL language that is not otherwise in error as far as that implementation is concerned.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

NOTE 15 — If a system is processing SQL language that contains errors, then it may be very difficult within a single statement to determine what is an error and what is an extension. As one possibility, an implementation may choose to check SQL language in two steps; first through its normal syntax analyzer and secondly through the SQL Flagger. The first step produces error messages for non-standard SQL language that the implementation cannot process or recognize. The second step processes SQL language that contains no errors as far as that implementation is concerned; it detects and flags at one time all non-standard SQL language that could be processed by that implementation. Any such two-step process should be transparent to the end user.

The SQL Flagger assists identification of conforming SQL language that may perform differently in alternative processing environments provided by a conforming SQL-implementation. It also provides a tool in identifying SQL elements that may have to be modified if SQL language is to be moved from a non-conforming to a conforming SQL processing environment.

An SQL Flagger provides one or more of the following "level of flagging" options:

— Core SQL Flagging

— Part SQL Flagging

— Package SQL Flagging

An SQL Flagger that provides one of these options shall be able to identify SQL language constructs that violate the indicated subset of SQL language. The subset of SQL language used by "Core SQL Flagging" is Core SQL. The subset of SQL language used by "Part SQL Flagging" is Core SQL plus those features required for conformance to a specified Part or Parts of ISO/IEC 9075. The subset of SQL language used by "Package Flagging" is Core SQL plus those features required for conformance to a specified Package or Packages of ISO/IEC 9075.

An SQL Flagger may also provide "SQL Feature Flagging". "SQL Feature Flagging" indicates which optional features in addition to the indicated subset of SQL language are needed to make the SQL language constructs conforming.

An SQL Flagger provides one or more of the following "extent of checking" options:

— Syntax Only

— Catalog Lookup

Under the Syntax Only option, the SQL Flagger analyzes only the SQL language that is presented; it checks for violations of any Syntax Rules that can be determined without access to the Information Schema. It does not necessarily detect violations that depend on the data type of syntactic elements, even if such violations are in principle deducible from the syntax alone.

NOTE 16 — Further details can be found in Annex F, "SQL feature taxonomy", in [ISO9075-2], Annex F, "SQL feature taxonomy", in [ISO9075-3], Annex F, "SQL feature taxonomy", in [ISO9075-4], Annex F, "SQL feature taxonomy", in [ISO9075-9], Annex F, "SQL feature taxonomy", in [ISO9075-10], Annex F, "SQL feature taxonomy", in [ISO9075-11], Annex F, "SQL feature taxonomy", in [ISO9075-13], and Annex F, "SQL feature taxonomy", in [ISO9075-14].

Under the Catalog Lookup option, the SQL Flagger assumes the availability of Definition Schema information and checks for violations of all Syntax Rules. For example, some Syntax Rules place restrictions on data types; this flagger option would identify extensions that relax such restrictions. In order to avoid security breaches, this option shall view the Definition Schema only through the eyes of a specific Information Schema. The flagger does not necessarily execute or simulate the execution of any <schema definition statement> or <schema manipulation statement>.



 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8.6 Claims of conformance

A claim of conformance to ISO/IEC 9075 shall include all of the following:

1) A claim of minimum conformance.

2) Zero or more claims of conformance to incremental parts.

3) Zero or more claims of conformance to packages.

4) Zero or more claims of conformance to optional features.

NOTE 17 — Each part of ISO/IEC 9075 specifies what shall be stated by claims of conformance to that part, in addition to the requirements of this clause.

8.6.1 Requirements for SQL applications

The term "SQL application" is used here to mean a collection of compilation units that contains one or more of:

— SQL statements.

— Invocations of SQL/CLI routines.

— Invocations of externally-invoked procedures.

A conforming SQL application shall be processed without syntax error, provided that all of the following is satisfied:

— Every SQL statement or SQL invocation is syntactically correct in accordance with ISO/IEC 9075.

— The schema contents satisfy the requirements of the SQL application.

— The SQL-data conforms to the schema contents.

— The user has not submitted for immediate execution an SQL-statement that is not syntactically correct.

A conforming SQL application shall not use any additional features, or features beyond the level of conformance claimed.

A claim of conformance by an SQL application shall also state:

— What implementation-defined elements and actions are relied on for correct performance.

— What schema contents are required to be supplied by the user.

8.6.2 Requirements for SQL-implementations

A conforming SQL-implementation shall process conforming SQL language according to the associated General Rules, Definitions, and Descriptions.

A conforming SQL implementation shall process conforming routine invocations according to the associated Definitions and General Rules.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8.6 Claims of conformance

A claim of conformance by an SQL-implementation shall also state:

— The definition for every element and action, within the scope of the claim, that ISO/IEC 9075 specifies to be implementation-defined.

A conforming SQL-implementation that provides additional facilities or that provides facilities beyond those specified as “Core” shall provide an SQL-Flagger.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex A
(informative)

Maintenance and interpretation of SQL

ISO/IEC JTC1 provides formal procedures for revision, maintenance, and interpretation of JTC1 Standards. Clause 14 of the JTC1 Directives, “Maintenance of International Standards”, specifies procedures for creating and processing “defect reports”. Defect reports may result in technical corrigenda, amendments, interpretations, or other commentary on an existing International Standard.

Potential new questions or new defect reports addressing the specifications of ISO/IEC 9075 should be communicated to:

Secretariat, ISO/IEC JTC1/SC32

American National Standards Institute

11 West 42nd Street

New York, NT 10036

USA

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex B
(informative)

Implementation-defined elements

This Annex references those features that are identified in the body of this part of ISO/IEC 9075 as implementation-defined.

The term *implementation-defined* is used to identify characteristics that may differ between implementations, but that shall ve defined for each particular implementation.

1) Subclause 4.2.2, "SQL-agents":

   a) In direct invocation of SQL (see Subclause 5.3.3.3, "Direct invocation of SQL"), the SQL-agent that causes the execution of SQL-statements is implementation-defined.

2) Subclause 4.2.4, "SQL-client modules":

   a) The mechanism by which an SQL-client module is created or destroyed is implementation-defined.

   b) If an SQL-client module has a name, its permitted names are implementation-defined.

3) Subclause 4.2.9.1, "Catalogs":

   a) The mechanisms for creating and destroying catalogs are implementation-defined.

   b) The default catalog name for execution of <preparable statement>s that are dynamically prepared in the current SQL-session through the execution of <prepare statement>s and <execute immediate statement>s are implementation-defined.

4) Subclause 4.2.9.2, "SQL-schemas":

   a) SQL-schemas may be created and destroyed by execution of SQL-schema statements or by mechanisms that are implementation-defined.

5) Subclause 4.4.3.1, "Numeric types":

   a) If the result of an arithmetic operation cannot be represented exactly in its result type, then whether the result is rounded or truncated is implementation-defined.

6) Subclause 4.4.3.2, "Character string types":

   a) The maximum length of a character string type is implementation-defined.

   b) The maximum variable length of a character large object type is implementation-defined.

7) Subclause 4.4.3.3, "Binary string types":

   a) The maximum length of a binary string type is implementation-defined.

   b) The maximum variable length of a binary large object type is implementation-defined.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8) Subclause 4.6.2.2, “Collations”:

   a) Supported collations not derived from a collation defined by a National or International Standard are implementation-defined.

9) Subclause 4.8.1, “Host languages”:

   a) The set of host languages supported by an SQL-implementation is implementation-defined.

   b) When using the SQL-client module binding style, the mechanism by which the SQL-client module is specified is implementation-defined.

   c) No matter what binding style is chosen, SQL-statements are written in an implementation-defined character set, known as the *source language character set*.

10) Subclause 4.8.2.1, “General parameter passing and data type correspondence information”:

   a) The method and time of binding between an external routine's reference to a compilation unit and that compilation unit is implementation-defined.

11) Subclause 4.10.1, “General routine information”:

   a) The method and time of binding between an external routine's reference to a compilation unit and that compilation unit is implementation-defined.

12) Subclause 4.11.1, “Classes of SQL-statements”:

   a) The additional statements provided by an SQL-implementation are implementation-defined.

13) Subclause 5.3.3.3, “Direct invocation of SQL”:

   a) In direct invocation of SQL, the methods of invoking SQL-statements, raising conditions, accessing diagnostic information, and reporting the results are implementation-defined.

14) Subclause 6.3.3.3, “Rule evaluation order”:

   a) Invocation of an SQL-invoked function is inessential if it is deterministic and does not possibly modify SQL-data; otherwise, it is implementation-defined whether it is essential or inessential.

15) Subclause 6.3.6, “Subclauses used as subroutines”:

   a) In a few cases, these Subclauses not having explicit syntax are intended to be invoked by other standards and/or through the use of implementation-defined mechanisms.

   b) The rules of these Subclauses behave as though they were a sort of definitional "subroutine" that is invoked by other Subclauses, other standards, or implementation-defined mechanisms.

16) Subclause 8.4.8, “Extensions and options”:

   a) An SQL-implementation may provide implementation-defined features that are additional to those specified by ISO/IEC 9075, and may add to the list of reserved words.

   b) It is implementation-defined whether an SQL Flagger flags implementation-defined features.

   c) An SQL-implementation may provide support for additional implementation-defined SQL-invoked routines or for implementation-defined argument values for SQL-invoked routines.

 © ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex C
(informative)

Implementation-dependent elements

This Annex references those places where this part of ISO/IEC 9075 states explicitly that the actions of a conforming implementation are implementation-dependent.

The term *implementation-dependent* is used to identify characteristics that may differ between implementations, but that are not necessarily specified for any particular implementation.

1) Subclause 4.6.1, "General SQL-schema object information":

   a) The specific name of an SQL-invoked routine, if not specified explicitly, is implementation-dependent.

2) Subclause 4.6.2.1, "Character sets":

   a) For characters contained entirely within an SQL-implementation, the methods for encoding them and collecting them into strings are implementation-dependent.

3) Subclause 4.7.2, "Determinism and constraints":

   a) The effect of invoking a routine that claims to be deterministic, but that is not in fact deterministic, is implementation-dependent.

4) Subclause 4.8.1, "Host languages":

   a) For the binding style Embedded SQL, the mechanisms for generating externally-invoked procedures from SQL-statements, for collecting those externally-invoked procedures into SQL-client modules, and for replacing each SQL-statement with an invocation of the externally-invoked procedure generated from it, are implementation-dependent.

5) Subclause 6.3.2, "Specification of the Information and Definition Schemata":

   a) The actual objects upon which Information Schema views are based are implementation-dependent.

6) Subclause 6.3.3.2, "Terms denoting rule requirements":

   a) The treatment of language that does not conform to the Formats and Syntax Rules of ISO/IEC 9075 is implementation-dependent.

7) Subclause 6.3.3.7, "Exceptions":

   a) The effect on any assignment targets and SQL descriptor areas of an SQL-statement that terminates with an exception condition, unless explicitly defined by ISO/IEC 9075, is implementation-dependent.

   b) If more than one condition could have occurred as a result of execution of an SQL-statement, it is implementation-dependent whether diagnostic information pertaining to more than one such condition is made available.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

8) Subclause 6.3.3.3, “Rule evaluation order”:

a) It is implementation-dependent whether expressions are actually evaluated from left to right when the precedence is not otherwise determined by the Formats or by parentheses.

b) If evaluation of the inessential parts of an expression or search condition would cause an exception condition to be raised, then it is implementation-dependent whether or not that condition is raised.

c) The declared type of a site that contains an intermediate result is implementation-dependent.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex D
(infomative)

Deprecated features

This Annex modifies Annex D, "Deprecated features", in ISO/IEC 9075-2.

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.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex E
(informative)

Incompatibilities with ISO/IEC 9075:2003

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 9075-1:2003.

Except as specified in this Annex, features and capabilities of Database Language SQL are compatible with ISO/IEC 9075-1:2003.

1) In ISO/IEC 9075-1:2003, an (Open Systems Interconnect, or OSI) Object Identifier was defined for Database Language SQL. In this edition of this part of ISO/IEC 9075, the definition of that Object Identifier has been eliminated.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

--`,,```,,,,````-`-`,,`,,`,`,,`---
© ISO/IEC 2008 – All rights reserved
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex F
(informative)

SQL feature taxonomy

This Annex describes a taxonomy of features and packages defined in this part of ISO/IEC 9075.

Table 4, "Feature taxonomy for optional features", contains a taxonomy of the optional features and packages 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 "Package ID" column of Table 4, "Feature taxonomy for optional features", specifies the formal identification of each feature and each subfeature contained in the table.

The "Package Name" column of Table 4, "Feature taxonomy for optional features", contains a brief description of the feature or subfeature associated with the Feature ID value.

Table 4, "Feature taxonomy for optional features", does not provide definitions of the features; the definition of those features is found in the Conformance Rules of the various parts of ISO/IEC 9075.

Table 4 — Feature taxonomy for optional features

	Package ID	Package Name
1	**PKG001**	**Enhanced datetime facilities**
2	**PKG002**	**Enhanced integrity management**
3	**PKG004**	**PSM**
4	**PKG006**	**Basic object support**
5	**PKG007**	**Enhanced object support**
6	**PKG008**	**Active database**
7	**PKG010**	**OLAP**

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Annex G
(informative)

Defect reports not addressed in this edition of this part 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.

None.

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS Not for Resale

(Blank page)

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS

Not for Resale

Bibliography

[1] [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*.

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

(Blank page)

--`,,```,,,,````-`-`,,`,,`,`,,`---
© ISO/IEC 2008 – All rights reserved
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

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> • 40, **41**
access mode • *27*
AFTER • 22
approximate numeric • *15*
array • *17*
array • 17
assertion • *22*
assertion • 44
assignment • *18*
assignment • 34, 43, 67
atomic • **5**
attribute • 6, 20, 43
attribute • *21*
attribute type • *20*
attribute values • *20*

— B —

base table • 17
base table • *21*
base tables • 9, 12, 13, 17, 18, 19, 20, 21, 35
base tables • *13*
be included in • *44*
BEFORE • 22
binary large object string • *15*
binary large object type • *15*
binary large object type • 65
binary string type • *15*
binary string type • 15, 65
binding styles • *24*
BNF nonterminal symbol • *37*
BNF nonterminal symbol • 47
Boolean • *15*
Boolean • 32

— C —

CALL • 33
catalog • 5, 12, 65
catalog • *11*
character encoding form • *19*
character large object type • *15*
character large object type • 65
character repertoire • 15, 19
character repertoire • *19*
character set • 15, 19, 25, 38, 66
character set • *19*
character string type • *15*
character string type • 15, 19, 65
check constraint • *21*
check constraint • 22
claim of minimum conformance • *55*
claim of minimum conformance • 55, 61
collation • *19*
collation • 19, 66
collection • 6, 11, 12, 13, 17, 43, 44, 61
collection • *17*
column • 6, 13, 20, 21, 24, 44, 45, 47, 56, 73
column • *21*
compilation unit • **5**
conditions • *26*
conditions • 26, 33, 34, 41, 66
conforming SQL-implementation • 5, 24, 26, 56, 59, 60, 61, 62
conforming SQL-implementation • *6*
constitute • x, 31
constitute • *40*
constraint mode • *23*
constraint mode • 23, 24
constructed • *14*
contain • 9, 13, 18, 24, 33, 35, 40, 45, 48

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

contain • *40*
contained in • *40*
contained in • 42, 43, 56, 73
containing • 9, 27, 41, 42, 47, 49, 50
containing • *40*
Core SQL • *55*
Core SQL • 56, 60
created base table • *13*
current base value • *23*

— D —

data type • **5**
datetime types • *16*
day-time intervals • *16*
declared temporary table • *13*
default collation • *19*
deferrable • *23*
deferrable • 23
deferred • *23*
definition schema • *12*
depends on • 22, 44
depends on • *44*
derived table • *13*
descriptor • **5**
descriptor area • *26*
descriptor area • 26, 33, 34, 43
diagnostics area • 10, 26
diagnostics area • *26*
direct result • *43*
direct supertype • *20*
direct supertype • 20
directly contains • *40*
distinct type • *21*
distinct type • 21
domain • 18, 20, 23, 45
domain • *20*
domain constraint • *20*

— E —

effectively • *41*
effectively • 41, 45, 47, 56
element type • *17*
encapsulated • *20*
END • 22
essential • *42*
essential • 66
exact numeric • *15*
exact numeric • 23
external routine • 25, 66
external routine • *28*
external sequence generator • *23*
externally-invoked procedures • 9, 61
externally-invoked procedures • *11*

— F —

Feature F052, “Intervals and datetime arithmetic” • 57
Feature F191, “Referential delete actions” • 57
Feature F411, “Time zone specification” • 57
Feature F491, “Constraint management” • 57
Feature F521, “Assertions” • 57
Feature F555, “Enhanced seconds precision” • 57
Feature F671, “Subqueries in CHECK constraints” • 57
Feature F701, “Referential update actions” • 57
fully qualified • **5**

— G —

generally contain • *40*
generally depends on • *44*
generally includes • *44*
global temporary table • *14*

— H —

host language • 11, 24, 25, 29, 33
host language • *24*, *32*

— I —

identifier • **5**
identify • **5**
immediate • *23*
immediate • 24, 61
immediately contain • *40*
implementation-defined • x, 6, 9, 10, 12, 15, 19, 24, 25, 28, 33, 42, 50, 59, 61, 62, 65, 66
implementation-defined • **5**
implementation-defined • *65*
implementation-defined features • *50*
implementation-defined features • 59, 66
implementation-dependent • **5**
implementation-dependent • 18, 19, 24, 39, 41, 42, 43, 44, 50, 67, 68
implementation-dependent • *67*
include • 12, 16, 34, 35, 36, 48, 55, 61
include • *44*
incremental parts • *45*
incremental parts • 55, 59, 61
indicator parameter • *26*

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

indicator parameter • 26
indirect result • *43*
inessential • *42*
inessential • 42, 66, 68
inherits • *20*
innermost • 38
innermost • *40*
instance • **5**
INSTEAD • 22
intermediate results • *42*
internal sequence generator • *23*
interval type • *16*
is dependent on • *44*
is generally dependent on • *44*
is part of • 40, 46
is part of • *43*
isolation level • *27*
isolation level • 27

— K —

known not nullable • *18*
known not nullable • 18

— L —

large object character string • *15*
local temporary table • *14*
locator • *20*, *25*
locator • 25

— M —

match type • *22*
methods • 19, 24, 28, 31, 32, 36, 66, 67
methods • *28*
<modified item> • **47**

— N —

national character set • *15*
not deferrable • *23*
null value • **5**
null value • *14*
null value • 14, 15, 18, 21, 26
nullability characteristic • *18*
nullability characteristic • 21

— O —

object • **5**
observer function • *20*
OF • 22
outermost • *40*
outermost • 43
overloaded • *28*

— P —

Feature P001, “Stored modules” • 58
Feature P002, “Computational completeness” • 58
Feature P003, “Information Schema views” • 58
persistent • **5**
persistent • 12, 13, 17, 18, 23, 27, 29, 31, 35
persistent • *17*
persistent base table • *13*
possibly non-deterministic • *24*
possibly non-deterministic • 24
possibly nullable • *18*
precision • 14, 15, 16, 18
precision • *15*
predefined • *14*
predefined • 14, 21, 36
primary key constraint • *21*
primary key constraint • 21
privilege • *23*
privilege • 23
production rule • *37*
production rule • 37, 38, 40
property • **6**

— Q —

query • *13*
query • 13, 22, 58

— R —

reference type • 6
reference type • *16*
referenced columns • *21*
referenced columns • 21
referenced table • *21*
referencing columns • *21*
referencing columns • 21
Referential actions • *22*
referential constraint • *21*
result parameter • *28*
result parameter • 28
RETURN • 33
role • *11*
role • 11, 23
role authorization • *11*

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

role authorization • 11
row • **6**

— S —

Feature S023, “Basic structured types” • 58
Feature S024, “Enhanced structured types” • 58
Feature S041, “Basic reference types” • 58
Feature S043, “Enhanced reference types” • 58
Feature S051, “Create table of type” • 58
Feature S071, “SQL-paths in function and type name resolution” • 58
Feature S081, “Subtables” • 58
Feature S111, “ONLY in query expressions” • 58
Feature S151, “Type predicate” • 58
Feature S161, “Subtype treatment” • 58
Feature S211, “User-defined cast functions” • 58
Feature S231, “Structured type locators” • 58
Feature S241, “Transform functions” • 58
schema • *12*
schema routine • *22*
scope • **6**
search condition • *22*
search condition • 22, 42
sequence • **6**
sequence generator • *23*
sequence generator • 23
serializable • *27*
shall • 19, 22, 23, 29, 37, 39, 42, 46, 55, 56, 59, 60, 61, 62
shall • *41*
simply contained in • *40*
simply containing • *40*
simply contains • *40*
simply contains • 40
site • 5, 16, 17, 18, 19, 21, 42, 43, 68
site • **6**
site • *17*
sizing item • **6**
source • *18*
source • 25, 35, 38
source language character set • *25*, *66*
specified by • 5, 6, 14, 15, 17, 19, 20, 24, 41, 45, 53, 59, 66
specified by • *40*
specify • *40*
SQL routine • *28*
SQL routine • 35
SQL-agent • 6, 9, 10, 11, 24, 25, 26, 27, 31, 32, 33, 34, 65
SQL-agent • *9*
SQL-client • 6, 9, 10, 11, 12, 14, 24, 25, 27, 29, 33, 34, 65, 66, 67
SQL-client • *10*
SQL-client module • *10*
SQL-connection • **6**
SQL-data • x, x, 6, 9, 10, 12, 13, 18, 21, 25, 26, 27, 29, 33, 41, 42, 43, 61, 66
SQL-data • *13*
SQL-environment • **6**
SQL-implementation • 1, 5, 6, 9, 10, 12, 13, 14, 17, 19, 24, 26, 31, 32, 34, 35, 53, 55, 56, 59, 62, 66, 67
SQL-implementation • **6**
SQL-implementation • *10*
SQL-invoked function • *28*
SQL-invoked function • 28, 42, 66
SQL-invoked procedure • *28*
SQL-invoked routine • 11, 22, 28, 43, 67
SQL-invoked routine • *28*
SQL-schema • 5, 12, 18, 35, 65, 67
SQL-schema • *12*
SQL-schema object • *12*
SQL-server • 6, 10, 12, 13, 26, 27, 29, 33, 34, 35
SQL-server • *10*
SQL-server module • *22*
SQL-session • **6**
SQL-statement • **6**
SQL-transaction • 5, 14, 17, 23, 26, 27, 29
SQL-transaction • *26*
standard-defined features • *50*
status parameter • *25*
status parameter • 44
structured type • 13, 16, 20, 21
structured type • *20*
subject table • *22*
substitutability • *20*
subtable • *20*, *21*
subtype • 6, 13, 20
subtype • *20*
supertables • *21*
syntax error or access rule violation • 41, 42

— T —

Feature T041, “Basic LOB data type support” • 58
Feature T191, “Referential action RESTRICT” • 57
Feature T201, “Comparable data types for referential constraints” • 57
Feature T211, “Basic trigger capability” • 57, 58
Feature T212, “Enhanced trigger capability” • 57

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

Feature T322, “Overloading of SQL-invoked functions and SQL-invoked procedures” • 57
Feature T431, “Extended grouping capabilities” • 59
Feature T611, “Elementary OLAP operators” • 59
table • 6, 13, 14, 17, 18, 20, 21, 22, 23, 29, 43, 44, 45, 47, 56, 58, 73
table • **6**
table • *13*
table constraint • *21*
target • *18*
target • 43
temporary • 13, 14, 17, 29
temporary • *17*
temporary table • *14*
transaction • *26*
transaction state • *27*
transaction state • 27
transaction-initiating statement • *26*
transient • *17*
transliteration • *19*
trigger • *22*
trigger • 22, 43, 57, 58
trigger action time • *22*
trigger event • *22*
trigger event • 22
triggered actions • *22*
triggered actions • 22
type preserving function • *28*
type preserving function • 28

— U —

unique columns • *21*
unique columns • 21
unique constraint • *21*
unique constraint • 21
updatable • *13*
user identifier • *11*
user-defined • *14*
user-defined • 14, 18, 20, 21, 28, 36, 43

— V —

view • 12, 20, 39, 40, 59, 60
view • *13*, *22*
view definition • *22*
viewed table • *13*
viewed table • 13, 22

— W —

with an intervening • *40*
with grant option • *23*
with grant option • 23
with hierarchy option • *23*
without an intervening • *40*
without an intervening • 40

— Y —

year-month intervals • *16*

© ISO/IEC 2008 – All rights reserved

Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale

ICS 35.060

Price based on 72 pages

© ISO/IEC 2008 – All rights reserved
Copyright International Organization for Standardization
Provided by IHS under license with ISO
No reproduction or networking permitted without license from IHS
Not for Resale