
Information technology — Database languages — SQL —

Part 1: Framework (SQL/Framework)

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

Partie 1: Charpente (SQL/Charpente)



Reference number
ISO/IEC 9075-1:2023(E)

AVA Documents
Document Sharing Hub
© ISO/IEC 2023



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2023

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

Foreword.....	x
Introduction.....	xii
1 Scope.....	1
2 Normative references.....	2
3 Terms and definitions.....	3
3.1 Definitions taken from ISO/IEC 10646:2020	3
3.2 Definitions provided in this document.....	3
4 Concepts.....	6
4.1 What is SQL?.....	6
4.2 Use of terms.....	6
4.3 Caveat.....	7
4.4 SQL-environments and their components.....	7
4.4.1 SQL-environments.....	7
4.4.2 SQL-agents.....	7
4.4.3 SQL-implementations.....	7
4.4.3.1 Introduction to SQL-implementations.....	7
4.4.3.2 SQL-clients.....	7
4.4.3.3 SQL-servers.....	8
4.4.4 SQL-client modules.....	8
4.4.5 User identifiers.....	8
4.4.6 Roles.....	8
4.4.7 User mapping concepts.....	8
4.4.8 Routine mapping concepts.....	8
4.4.9 Catalogs and schemas.....	9
4.4.9.1 Catalogs.....	9
4.4.9.2 SQL-schemas.....	9
4.4.9.3 Information Schema.....	9
4.4.9.4 Definition Schema.....	9
4.4.10 Foreign servers and descriptors.....	9
4.4.11 Foreign-data wrappers and descriptors.....	9
4.4.12 SQL-data.....	10
4.5 Tables.....	10
4.6 SQL data types.....	11
4.6.1 General data type information.....	11
4.6.2 Null value.....	11
4.6.3 Predefined types.....	11
4.6.3.1 Numeric types.....	11
4.6.3.2 Character string types.....	12

4.6.3.3	Binary string types.	12
4.6.3.4	Boolean type.	12
4.6.3.5	Datetime types.	12
4.6.3.6	Interval types.	12
4.6.3.7	XML type.	12
4.6.3.8	JSON type.	13
4.6.4	Constructed atomic types: reference types.	13
4.6.5	Constructed composite types.	13
4.6.5.1	Collection types.	13
4.6.5.2	Row types.	13
4.7	Sites and operations on sites.	13
4.7.1	Sites.	13
4.7.2	Assignment.	13
4.7.3	Nullability.	14
4.8	SQL-schema objects.	14
4.8.1	General SQL-schema object information.	14
4.8.2	Descriptors relating to character sets.	14
4.8.2.1	Character sets.	14
4.8.2.2	Collations.	15
4.8.2.3	Transliterations.	15
4.8.3	Domains and their components.	15
4.8.3.1	Domains.	15
4.8.3.2	Domain constraints.	15
4.8.4	User-defined types.	15
4.8.4.1	Introduction to user-defined types.	15
4.8.4.2	Distinct types.	15
4.8.4.3	Structured types.	16
4.8.5	Base tables and their components.	16
4.8.5.1	Base tables.	16
4.8.5.2	Columns.	16
4.8.5.3	Periods.	16
4.8.5.4	Table constraints.	16
4.8.5.5	Triggers.	17
4.8.6	View definitions.	17
4.8.7	Assertions.	17
4.8.8	SQL-server modules (defined in ISO/IEC 9075-4).	17
4.8.9	Schema routines.	17
4.8.10	Sequence generators.	17
4.8.11	Privileges.	17
4.9	Integrity constraints and constraint checking.	18
4.9.1	Constraint checking.	18
4.9.2	Determinism and constraints.	18
4.9.3	Consistency when deleting and updating multiple rows.	19
4.10	Communication between an SQL-agent and an SQL-server.	19
4.10.1	Host languages.	19
4.10.2	Source language character set.	20
4.10.3	Parameter passing and data type correspondences.	20

4.10.3.1	General parameter passing and data type correspondence information.	20
4.10.3.2	Data type correspondences.	20
4.10.3.3	Locators.	21
4.10.3.4	Status parameters.	21
4.10.3.5	Indicator parameters.	21
4.10.4	Descriptor areas.	21
4.10.5	Diagnostic information.	21
4.10.6	SQL-transactions.	21
4.11	Modules.	22
4.12	Routines.	22
4.12.1	General routine information.	22
4.12.2	Type preserving functions.	23
4.13	SQL-statements.	23
4.13.1	Classes of SQL-statements.	23
4.13.2	SQL-statements classified by function.	24
5	Parts of the ISO/IEC 9075 series.	25
5.1	Overview.	25
5.2	ISO/IEC 9075-1: Framework (SQL/Framework).	25
5.3	ISO/IEC 9075-2 Foundation (SQL/Foundation).	26
5.3.1	Introduction to ISO/IEC 9075-2 Foundation (SQL/Foundation).	26
5.3.2	Bindings methods.	26
5.3.2.1	Introduction to bindings methods.	26
5.3.2.2	Embedded SQL.	26
5.3.2.3	Dynamic SQL.	26
5.3.2.4	Direct invocation of SQL.	26
5.3.3	SQL-statements specified in ISO/IEC 9075-2	26
5.4	ISO/IEC 9075-3 : Call-Level Interface (SQL/CLI).	27
5.5	ISO/IEC 9075-4 : Persistent Stored Modules (SQL/PSM).	28
5.5.1	Introduction to SQL/PSM.	28
5.5.2	SQL-statements specified in ISO/IEC 9075-4	28
5.6	ISO/IEC 9075-9 : Management of External Data (SQL/MED).	28
5.7	ISO/IEC 9075-10 : Object Language Bindings (SQL/OLB).	28
5.8	ISO/IEC 9075-11 : Information and Definition Schemas (SQL/Schemata).	28
5.9	ISO/IEC 9075-13 : SQL Routines and Types Using the Java™ Programming Language (SQL/JRT).	29
5.10	ISO/IEC 9075-14 : XML-Related Specifications (SQL/XML).	29
5.11	ISO/IEC 9075-15 : Multidimensional Arrays (SQL/MDA).	29
5.12	ISO/IEC 9075-16 : Property Graph Queries (SQL/PGQ).	29
6	Notation and conventions used in other parts of the ISO/IEC 9075 series.	31
6.1	Notation taken from ISO/IEC 10646:2020	31
6.2	Notation provided in the ISO/IEC 9075 series.	31
6.3	Conventions.	32
6.3.1	Specification of syntactic elements.	32
6.3.2	Specification of the Information and Definition Schemata.	33
6.3.3	Use of terms.	33
6.3.3.1	Syntactic containment.	33
6.3.3.2	Terms denoting rule requirements.	34
6.3.3.3	Rule evaluation order.	35

6.3.3.4	Conditional rules.	36
6.3.3.5	Syntactic substitution.	36
6.3.3.6	Other terms.	37
6.3.3.7	Exceptions.	37
6.3.4	Descriptors.	38
6.3.5	Application of technical corrigenda.	39
6.3.6	Relationships of parts within the ISO/IEC 9075 series.	39
6.3.6.1	Introduction to relationships among parts.	39
6.3.6.2	Clauses 1, 2, and 3.	40
6.3.6.3	New and modified Clauses and Subclauses.	40
6.3.6.4	New and modified Tables, Figures, Examples, and Equations.	41
6.3.6.5	Functions.	41
6.3.6.6	New and modified Format Items.	41
6.3.6.7	New and modified paragraphs and rules.	42
6.3.6.8	Modified Subclause Signatures.	44
6.3.6.9	New and modified Annexes.	44
6.3.6.10	Order of merging an incremental part.	45
6.3.7	Subclauses used as subroutines.	45
6.3.8	Document typography.	46
6.3.9	Index typography.	46
6.3.10	Feature ID and Feature Name.	46
6.3.11	Implementation-defined and implementation-dependent.	47
6.4	Digital artifacts.	48
6.4.1	Introduction to digital artifacts.	48
6.4.2	Language syntax.	48
6.4.3	Condition codes.	48
6.4.4	Feature codes.	49
6.4.5	Implementation-defined items.	50
6.4.6	Implementation-dependent items.	50
6.4.7	Header files.	51
6.4.8	Ada interface.	51
6.4.9	Schema definition.	51
6.4.10	XML Schemata.	52
7	Annexes to the parts of the ISO/IEC 9075 series.	53
7.1	Annexes are informative.	53
7.2	SQL conformance summary.	53
7.3	Implementation-defined elements.	53
7.4	Implementation-dependent elements.	53
7.5	Deprecated features.	53
7.6	Incompatibilities with previous versions.	53
7.7	SQL feature taxonomy.	53
7.8	Defect Reports.	53
8	Status codes.	54
8.1	SQLSTATE.	54
9	Conformance.	55
9.1	Kinds of conformance claims.	55

9.2	Minimum conformance.....	55
9.3	Conformance to parts.....	55
9.4	Conformance to features.....	55
9.5	Extensions and options.....	56
9.6	SQL flagger.....	56
9.7	Claims of conformance.....	57
9.7.1	How conformance is claimed.....	57
9.7.2	Requirements for SQL applications.....	58
9.7.3	Requirements for SQL-implementations.....	58
Annex A (informative) SQL conformance summary.....		59
Annex B (informative) Implementation-defined elements.....		60
Annex C (informative) Implementation-dependent elements.....		63
Annex D (informative) SQL optional feature taxonomy.....		65
Annex E (informative) Deprecated features.....		66
Annex F (informative) Incompatibilities with ISO/IEC 9075:2016.....		67
Annex G (informative) Defect Reports not addressed in this edition of this document.....		68
Annex H (informative) Maintenance and interpretation of SQL.....		69
Annex I (informative) Support for the use of inter-document links in the ISO/IEC 9075 series..		70
Bibliography.....		71
Index.....		72

Tables

Table	Page
1 Relationships between externally-invoked and SQL-invoked routines.	22
2 Symbols used in BNF.	31
3 Examples and explanations of paragraph and rule mergers.	43
4 SQLSTATE class and subclass codes.	54
A.1 Feature definitions outside of Conformance Rules.	59
D.1 Feature taxonomy for optional features.	65
I.1 Filename conventions.	70

Figures

Figure	Page
1 Relationships between the parts of the ISO/IEC 9075 series.	40

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.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

ISO and IEC draw attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO and IEC take no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO and IEC have not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents and <https://patents.iec.ch>. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see www.iso.org/iso/foreword.html. In the IEC, see www.iec.ch/understanding-standards.

This document was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 32, *Data management and interchange*.

This sixth edition cancels and replaces the fifth edition (ISO/IEC 9075-1:2016), which has been technically revised. It also incorporates the Technical Corrigendum ISO/IEC 9075-1:2016/Cor.1:2022.

The main changes are as follows:

- addition and refinement of the terms and concepts to support new data types required by incremental parts;
- clarification and correction of the merge instructions for Technical Corrigenda and incremental parts;
- improve the presentation and accuracy of the summaries of implementation-defined and implementation-dependent aspects of this document;
- introduction of several digital artifacts;
- alignment with updated ISO house style and other guidelines for creating standards.

This sixth edition of ISO/IEC 9075-1 is designed to be used in conjunction with the following editions of other parts of the ISO/IEC 9075 series, all published 2023:

- ISO/IEC 9075-2, sixth edition;
- ISO/IEC 9075-3, sixth edition;
- ISO/IEC 9075-4, seventh edition;
- ISO/IEC 9075-9, fifth edition;
- ISO/IEC 9075-10, fifth edition;
- ISO/IEC 9075-11, fifth edition;
- ISO/IEC 9075-13, fifth edition;
- ISO/IEC 9075-14, sixth edition;
- ISO/IEC 9075-15, second edition;
- ISO/IEC 9075-16, first edition.

A list of all parts in the ISO/IEC 9075 series can be found on the ISO and IEC websites.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html and www.iec.ch/national-committees.

Introduction

The organization of this document is as follows:

- 1) [Clause 1, “Scope”](#), specifies the scope of this document.
- 2) [Clause 2, “Normative references”](#), identifies additional standards that, through reference in this document, constitute provisions of the ISO/IEC 9075 series.
- 3) [Clause 3, “Terms and definitions”](#), defines terms and definitions used in this document and in other parts of the ISO/IEC 9075 series.
- 4) [Clause 4, “Concepts”](#), describes the concepts used in the ISO/IEC 9075 series.
- 5) [Clause 5, “Parts of the ISO/IEC 9075 series”](#), summarizes the content of each of the parts of the ISO/IEC 9075 series, in terms of the concepts described in [Clause 4, “Concepts”](#).
- 6) [Clause 6, “Notation and conventions used in other parts of the ISO/IEC 9075 series”](#), defines notation and conventions used in other parts of the ISO/IEC 9075 series.
- 7) [Clause 7, “Annexes to the parts of the ISO/IEC 9075 series”](#), describes the content of annexes of other parts of the ISO/IEC 9075 series.
- 8) [Clause 8, “Status codes”](#), defines values that identify the status of the execution of SQL-statements and the mechanisms by which those values are returned.
- 9) [Clause 9, “Conformance”](#), specifies requirements that apply to claims of conformance to all or some of the parts of the ISO/IEC 9075 series.
- 10) [Annex A, “SQL conformance summary”](#), is an informative Annex. It summarizes the conformance requirements of the SQL language.
- 11) [Annex B, “Implementation-defined elements”](#), is an informative Annex. It lists those features for which the body of this document states that the syntax, the meaning, the returned results, the effect on SQL-data and/or schemas, or other aspect is partly or wholly implementation-defined.
- 12) [Annex C, “Implementation-dependent elements”](#), is an informative Annex. It lists those features for which the body of this document states that the syntax, the meaning, the returned results, the effect on SQL-data and/or schemas, or other aspect is partly or wholly implementation-dependent.
- 13) [Annex D, “SQL optional feature taxonomy”](#), is an informative Annex. It identifies features of the SQL language specified in this document by an identifier and a short descriptive name. This taxonomy is used to specify conformance.
- 14) [Annex E, “Deprecated features”](#), is an informative Annex. It lists features that the responsible Technical Committee intends not to include in a future edition of this document.
- 15) [Annex F, “Incompatibilities with ISO/IEC 9075:2016”](#), is an informative Annex. It lists incompatibilities with the previous edition of this document.
- 16) [Annex G, “Defect Reports not addressed in this edition of this document”](#), is an informative Annex. It describes the Defect Reports that were known at the time of publication of this document. Each of these problems is a problem carried forward from the previous edition of this document. No new problems have been created in the drafting of this edition of this document.
- 17) [Annex H, “Maintenance and interpretation of SQL”](#), is an informative Annex. It describes the formal procedures for maintenance and interpretation of the ISO/IEC 9075 series.

- 18) [Annex I, “Support for the use of inter-document links in the ISO/IEC 9075 series”](#), is an informative Annex. It describes conventions that users of the ISO/IEC 9075 series use in order to maximize the facilities provided in the parts of the ISO/IEC 9075 series.

Information technology — Database language SQL —

Part 1:

Framework (SQL/Framework)

1 Scope

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

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

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements 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.

ISO/IEC 646, *Information technology — ISO 7-bit coded character set for information interchange*

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

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

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

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

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

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

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

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

ISO/IEC 9075-15, *Information technology — Database languages — SQL — Part 15: Multidimensional Arrays (SQL/MDA)*

ISO/IEC 9075-16, *Information technology — Database languages — SQL — Part 16: Property Graph Queries (SQL/PGQ)*

ISO/IEC 10646:2020, *Information technology — Universal Multi-Octet Coded Character Set (UCS)*

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

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- ISO Online browsing platform: available at <https://www.iso.org/obp>
- IEC Electropedia: available at <https://www.electropedia.org/>

3.1 Definitions taken from ISO/IEC 10646:2020

3.1.1

character

member of a set of elements used for the organization, control, or representation of textual data

Note 1 to entry: This is identical to *abstract character* as defined by [The Unicode® Standard](#). In this document, when the relevant character repertoire is UCS, a character can be thought of as *that which is represented by one code point*.

[SOURCE: ISO/IEC 10646:2020, 3.5]

3.2 Definitions provided in this document

3.2.1

atomic, adj.

incapable of being subdivided

3.2.2

column

field of the row type of a table

3.2.3

compilation unit

segment of executable code, possibly consisting of one or more subprograms

3.2.4

data type

set of representable values

3.2.5

descriptor

coded description of the metadata of an SQL object

3.2.6

field

(field name, data type) pair

Note 1 to entry: A value of a field is a value of its data type.

3.2.7

identifier

object by which something is identified

3.2.8

identify, verb

to reference something without ambiguity

3.2.9

implementation-defined, adj.

possibly differing between SQL-implementations, but specified by the implementor for each particular SQL-implementation

3.2.10

implementation-dependent, adj.

possibly differing between SQL-implementations, but not specified by the ISO/IEC 9075 series, and not required to be specified by the implementor for a particular SQL-implementation

3.2.11

instance

<of a value> physical representation of a value

Note 1 to entry: Each instance is at exactly one site. An instance has a data type that is the data type of its value.

3.2.12

null value

special value that is used to indicate the absence of any data value

3.2.13

object

<as in “x object”> *thing* that is a component of, or is otherwise associated with, some x and cannot exist independently of that x

Note 1 to entry: 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.

3.2.14

persistent, adj.

continuing to exist indefinitely, until destroyed deliberately

Note 1 to entry: 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.

3.2.15

property

<of an object> attribute, quality, or characteristic of the object

3.2.16

row

value of a row type, i.e., a sequence of field values; for rows in a table, a sequence of column values

3.2.17

scope

<of a standard> 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

3.2.18

scope

<of a name or declaration> one or more BNF non-terminal symbols within which the name or declaration is effective

3.2.19

scope

<of a reference type> table that a value of that reference type references

3.2.20

sequence

ordered collection of objects that are not necessarily distinct

3.2.21

site

place occupied by an instance of a value of some specified data type (or subtype of that data type)

3.2.22

SQL-connection

association between an SQL-client and an SQL-server

3.2.23

SQL-environment

context in which SQL-data exists and SQL-statements are executed

3.2.24

SQL-implementation

processor that processes SQL-statements

Note 1 to entry: A *conforming SQL-implementation* is an SQL-implementation that satisfies the requirements for SQL-implementations as defined in [Clause 9](#), “*Conformance*”.

3.2.25

SQL-session

context within which a single user, from a single SQL-agent, executes a sequence of consecutive SQL-statements over a single SQL-connection

3.2.26

SQL-statement

string of characters that conforms to the Format and Syntax Rules specified in the parts of the ISO/IEC 9075 series

3.2.27

table

collection of zero or more rows

3.2.28

whitespace

sequence of one or more characters that have the Unicode property `White_Space`

4 Concepts

4.1 What is SQL?

The name “SQL” is correctly pronounced “ess cue ell,” instead of the somewhat common “sequel”. SQL, specified in the parts of the ISO/IEC 9075 series, is a *database language* (more precisely, a *data sublanguage*) used for access to pseudo-relational databases that are managed by pseudo-relational database management systems (RDBMS).

NOTE 1 — Many books and articles “define” SQL by parenthetically claiming that the letters stand for “Structured Query Language”. While this was undoubtedly true for the original prototypes, it is not true of the standard. When the letters appear in product names, they have often been assigned this meaning by the product implementors, but users are ill-served by claims that the word “structured” accurately describes the language. In the SQL standard, the letters do not stand for anything at all. There are aficionados who assert that the letters stand for “SQL Query Language”, a recursive acronym in the tradition of GNU (GNU’s Not Unix) and SPARQL (SPARQL Protocol And RDF Query Language).

SQL is *based on*, but is not a strict implementation of, the relational model of data [Relational Model](#), making SQL “pseudo-relational” instead of truly relational. There are two principal ways in which SQL breaks from the relational model.

- The relational model requires that every *relation*¹ have no duplicate rows (that is, rows for which all corresponding attributes have equal values). A consequence of this requirement is that every relation has a primary key. SQL does not enforce this requirement, but provides syntax that permits application developers to enforce the requirement as needed.
- The relational model does not specify or recognize any sort of *flag* or other marker that represents unspecified, unknown, or otherwise missing data values. Consequently, the relational model depends only on two-valued (true/false) logic. SQL provides a “null value” that serves this purpose. In support of null values, SQL also depends on three-valued (true/false/unknown) logic.

The SQL standard defines two distinct aspects of the language: a *data definition language* (known in the standard as a “schema manipulation language” that supports the creation and destruction of database objects, such as tables, columns, and integrity constraints; and a *data manipulation language* that allows the retrieval, creation, modification, and deletion of data.

An especially important characteristic of SQL is that it is *declarative* (as opposed to *procedural*) in nature.² That is, SQL’s syntax and semantics allow application writers to specify the rules by which results may be computed by an SQL system, rather than specifying the algorithms by which those results must be computed. This contrasts with earlier database languages that required the details specification of step-by-step algorithms to retrieve, filter, and format the data required by applications.

4.2 Use of terms

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

Some objects are a component of an object on which they depend. If an object ceases to exist, then every object dependent on that objects also ceases to exist. The metadata representation of an object is known as a descriptor. See also [Subclause 6.3.4, “Descriptors”](#).

¹ Terms defined by the relational model are used correctly in this Subclause, but are not defined or explained in this document; readers are referred to the formal specification of the relational model.

² ISO/IEC 9075-4, specifies a number of SQL statements that provide a *procedural* capability through which procedures and functions can be written entirely in SQL. Procedural programs can also be written in ordinary “host” programming languages with embedded SQL statements.

4.3 Caveat

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

4.4 SQL-environments and their components

4.4.1 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 the ISO/IEC 9075 series, because it has different meanings in the more general context.

4.4.2 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 5.3.2.4, “Direct invocation of SQL”](#)), it is implementation-defined ([IW002](#)). Alternatively, it may consist of one or more compilation units that, when executed, invoke externally-invoked procedures in an SQL-client module.

4.4.3 SQL-implementations

4.4.3.1 Introduction to 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 is able to 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 implementor and those determined by the installer.

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

4.4.3.2 SQL-clients

An *SQL-client* is a processor, perceived by the SQL-agent to be 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.

4.4 SQL-environments and their components

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 4.10.5, “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.

4.4.3.3 SQL-servers

Each *SQL-server* is a processor, perceived by the SQL-agent to be 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.

4.4.4 SQL-client modules

An *SQL-client module* is a module that is explicitly created and dropped by implementation-defined [\(IW003\)](#) mechanisms.

An SQL-client module does not necessarily have a name; if it does, the permitted names are implementation-defined [\(IA094\)](#).

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 given time. However, in the case of either direct binding style or SQL/CLI, this may be a default SQL-client module whose existence is not apparent to the user.

4.4.5 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 the ISO/IEC 9075 series.

4.4.6 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 4.8.11, “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.

4.4.7 User mapping concepts

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

4.4.8 Routine mapping concepts

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

NOTE 2 — In this revision of the ISO/IEC 9075 series, the SQL-invoked routine that is paired with a foreign server descriptor is restricted to be an SQL-invoked regular function. It is possible that this restriction will not appear in future editions.

4.4.9 Catalogs and schemas

4.4.9.1 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 (IW005).

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 (ID006) but may be changed by the use of <set catalog statement>s.

4.4.9.2 SQL-schemas

An *SQL-schema*, often referred to simply as a *schema*, is a persistent, named collection of descriptors. An 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 (IW006) mechanisms).

4.4.9.3 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 be 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 4.4.9.4, "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.

4.4.9.4 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 in a catalog. The ISO/IEC 9075 series 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.

4.4.10 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.

4.4.11 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.

4.4.12 SQL-data

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

4.5 Tables

A *table* is a collection of zero or more rows where each row is an instance of the row type of the table. 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>. The columns of a table are the fields of the row type of the table. 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 4.8.5, “Base tables and their components”](#)) or a module object (its descriptor is in a module; see [Subclause 4.11, “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 persistent base table is either a *regular persistent base table* or a *system-versioned table*. System-versioned tables differ from regular persistent base tables in two major respects:

- system-versioned tables have two columns of datetime type whose values are generated automatically;
- the update and delete operations on system-versioned tables maintain those columns and have additional side effects.

A base table declared in a module is always a temporary table, and is called a *declared local temporary table*.

A *temporary table* is an SQL-session object that cannot be 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.

4.6 SQL data types

4.6.1 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 the ISO/IEC 9075 series 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 4.8.4, “User-defined types”](#).

A predefined data type is a data type specified by the ISO/IEC 9075 series, 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.

4.6.2 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.

4.6.3 Predefined types

4.6.3.1 Numeric types

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

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 ([IA002](#)). 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.

4.6.3.2 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 (IL001) 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 (IL002) 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 SQL-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.

4.6.3.3 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 (IL003) 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 (IL004) maximum that is probably greater than that of other binary strings.

4.6.3.4 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.

4.6.3.5 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.

4.6.3.6 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 an interval 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.

4.6.3.7 XML type

Values of the XML type are called XML values.

4.6.3.8 JSON type

Values of the JSON type are called SQL/JSON values.

4.6.4 Constructed atomic types: 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.

4.6.5 Constructed composite types

4.6.5.1 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.

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

An MD-array type is specified by an MD-array type constructor.

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

A multiset type is specified by a multiset type constructor.

4.6.5.2 Row types

A row type is a sequence of one or more fields. A value of a row type is the sequence of values of its fields.

4.7 Sites and operations on sites

4.7.1 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 be *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 4.8.5.4, "Table constraints"](#)).

Some sites may be referenced by their names — for example, base tables and SQL variables (see [ISO/IEC 9075-4](#)). Some sites may be referenced by a REF value. A site occupied by an element of an array may be referenced by its element number.

4.7.2 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.