

Australian Standard™

**Information technology—Database
languages—SQL**

**Part 13: SQL Routines and Types Using
the Java™ Programming Language
(SQL/JRT)**



This Australian Standard was prepared by Committee IT-027, Data Management and Interchange. It was approved on behalf of the Council of Standards Australia on 16 May 2005. This Standard was published on 16 June 2005.

The following are represented on Committee IT-027:

Association of Superannuation Funds of Australia
Australia Post
Australia Bankers Association
Australian Bureau of Statistics
Australian Computer Society
Australian Customs Service
Australian Electoral Commission
Australian Electric and Electronic Manufacturers Association
Australian Industry Group
Australian Information Industry Association
Australian Institute of Health and Welfare
Australian Taxation Office
Centrelink
Department of Defence
Department of Immigration, Multicultural and Indigenous Affairs
Health Insurance Commission
Data Management Association Australia

Keeping Standards up-to-date

Standards are living documents which reflect progress in science, technology and systems. To maintain their currency, all Standards are periodically reviewed, and new editions are published. Between editions, amendments may be issued. Standards may also be withdrawn. It is important that readers assure themselves they are using a current Standard, which should include any amendments which may have been published since the Standard was purchased.

Detailed information about Standards can be found by visiting the Standards Web Shop at www.standards.com.au and looking up the relevant Standard in the on-line catalogue.

Alternatively, the printed Catalogue provides information current at 1 January each year, and the monthly magazine, *The Global Standard*, has a full listing of revisions and amendments published each month.

Australian Standards™ and other products and services developed by Standards Australia are published and distributed under contract by SAI Global, which operates the Standards Web Shop.

We also welcome suggestions for improvement in our Standards, and especially encourage readers to notify us immediately of any apparent inaccuracies or ambiguities. Contact us via email at mail@standards.org.au, or write to the Chief Executive, Standards Australia, GPO Box 5420, Sydney, NSW 2001.

This Standard was issued in draft form for comment as DR 05150.

Australian Standard™

Information technology—Database languages—SQL

Part 13: SQL Routines and Types Using the Java™ Programming Language (SQL/JRT)

First published as AS ISO/IEC 9075.13—2005.

COPYRIGHT

© Standards Australia

All rights are reserved. No part of this work may be reproduced or copied in any form or by any means, electronic or mechanical, including photocopying, without the written permission of the publisher.

Published by Standards Australia GPO Box 5420, Sydney, NSW 2001, Australia

ISBN 0 7337 6771 0

PREFACE

This Standard was prepared by the Standards Australia Committee IT-027, Data Management and Interchange.

This Standard is identical with, and has been reproduced from ISO/IEC 9075-13:2003, *Information technology—Database languages—SQL—Part 13: SQL Routines and Types Using the Java™ Programming Language (SQL/JRT)*.

The objective of this Standard is to provide database designers, administrators and developers with the specifics to invoke static methods written in the Java(tm) programming language as SQL invoked routines and to use classes defined in the Java programming language as SQL structured user-defined types.

This Standard is Part 13 of AS 9075, *Information technology—Database languages—SQL*, which is published in parts as follows:

- 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 5: Host Language Bindings (SQL/Bindings)
- Part 9: Management of External Data (SQL/MED)
- Part 10: Object Language Bindings (SQL/OLB)
- Part 11: Information and Definition Schemas (SQL/Schemata)
- Part 13: SQL Routines and Types Using the Java™ Programming Language (SQL/JRT) (this Standard)
- Part 14: XML-Related Specifications (SQL/XML)

As this Standard is reproduced from an international standard, the following applies:

- (a) Its number appears on the cover and title page while the international standard number appears only on the cover.
- (b) In the source text ‘this part of ISO/IEC 9075’ should read ‘this Australian Standard’.
- (c) A full point substitutes for a comma when referring to a decimal marker.

References to International Standards should be replaced by references to Australian or Australian/New Zealand Standards, as follows:

<i>Reference to International Standard</i>		<i>Australian Standard</i>	
ISO/IEC		AS	
9075	Information technology—Database languages—SQL	9075	Information technology—Database languages—SQL
9075-1	Part 1: Framework (SQL/Framework)	9075.1	Part 1: Framework (SQL/Framework)
9075-2	Part 2: Foundation (SQL/Foundation)	9075.2	Part 2: Foundation (SQL/Foundation)
9075-10	Part 10: Object Language Bindings (SQL/OLB)	9075.10	Part 10: Object Language Bindings (SQL/OLB)
9075-11	Part 11: Information and Definition Schemas (SQL/Schemas)	9075.11	Part 11: Information and Definition Schemas (SQL/Schemas)

CONTENTS

Page

1	Scope.....	1
2	Normative references.....	3
2.1	JTC1 standards.....	3
2.2	Other international standards.....	3
3	Definitions, notations, and conventions.....	5
3.1	Definitions.....	5
3.1.1	Definitions taken from [Java].....	5
3.1.2	Definitions taken from [JVM].....	5
3.1.3	Definitions provided in Part 13.....	6
3.2	Conventions.....	7
3.2.1	Specification of built-in procedures.....	7
3.2.2	Specification of deployment descriptor files.....	7
4	Concepts.....	9
4.1	The Java programming language.....	9
4.2	SQL-invoked routines.....	10
4.2.1	Overview of SQL-invoked routines.....	10
4.2.2	Characteristics of SQL-invoked routines.....	11
4.3	Java class name resolution.....	12
4.4	SQL result sets.....	13
4.5	Parameter mapping.....	13
4.6	Unhandled Java exceptions.....	14
4.7	Data types.....	15
4.7.1	Non-predefined and non-SQL types.....	15
4.8	User-defined types.....	15
4.8.1	Introduction to user-defined types.....	15
4.8.2	User-defined type descriptor.....	16
4.8.3	User-defined type comparison and assignment.....	18
4.8.4	Accessing static fields.....	18
4.8.5	Converting objects between SQL and Java.....	19
4.8.5.1	SERIALIZABLE.....	19
4.8.5.2	SQLDATA.....	19
4.8.5.3	Developing for portability.....	20
4.9	Built-in procedures.....	20
4.10	Privileges.....	21

4.11	JARs.....	21
4.11.1	Deployment descriptor files.....	22
5	Lexical elements.....	23
5.1	<token> and <separator>.....	23
5.2	Names and identifiers.....	25
6	Scalar expressions.....	27
6.1	<method invocation>.....	27
6.2	<new specification>.....	28
7	Predicates.....	29
7.1	<comparison predicate>.....	29
8	Additional common elements.....	31
8.1	<Java parameter declaration list>.....	31
8.2	<SQL Java path>.....	32
8.3	<routine invocation>.....	34
8.4	<language clause>.....	43
8.5	Execution of array-returning functions.....	44
8.6	Java routine signature determination.....	51
9	Schema definition and manipulation.....	61
9.1	<drop schema statement>.....	61
9.2	<table definition>.....	63
9.3	<view definition>.....	64
9.4	<user-defined type definition>.....	65
9.5	<attribute definition>.....	69
9.6	<alter type statement>.....	73
9.7	<drop data type statement>.....	74
9.8	<SQL-invoked routine>.....	75
9.9	<alter routine statement>.....	79
9.10	<drop routine statement>.....	80
9.11	<user-defined ordering definition>.....	81
9.12	<drop user-defined ordering statement>.....	82
10	Access control.....	83
10.1	<grant privilege statement>.....	83
10.2	<privileges>.....	84
10.3	<revoke statement>.....	85
11	Built-in procedures.....	87
11.1	SQLJ.INSTALL_JAR procedure.....	87
11.2	SQLJ.REPLACE_JAR procedure.....	89
11.3	SQLJ.REMOVE_JAR procedure.....	91
11.4	SQLJ.ALTER_JAVA_PATH procedure.....	93
12	Java topics.....	95

12.1	Java facilities supported by this part of ISO/IEC 9075.	95
12.1.1	Package java.sql.	95
12.1.2	System properties.	95
12.2	Deployment descriptor files.	97
13	Information Schema.	101
13.1	JAR_JAR_USAGE view.	101
13.2	JARS view.	102
13.3	METHOD_SPECIFICATIONS view.	103
13.4	ROUTINE_JAR_USAGE view.	104
13.5	TYPE_JAR_USAGE view.	105
13.6	USER_DEFINED_TYPES view.	106
13.7	Short name views.	107
14	Definition Schema.	109
14.1	JAR_JAR_USAGE base table.	109
14.2	JARS base table.	111
14.3	METHOD_SPECIFICATIONS base table.	112
14.4	ROUTINE_JAR_USAGE base table.	114
14.5	ROUTINES base table.	115
14.6	TYPE_JAR_USAGE base table.	116
14.7	USAGE_PRIVILEGES base table.	117
14.8	USER_DEFINED_TYPES base table.	118
15	Status codes.	121
15.1	Class and subclass values for uncaught Java exceptions.	121
15.2	SQLSTATE.	122
16	Conformance.	123
16.1	Claims of conformance to SQL/JRT.	123
16.2	Additional conformance requirements for SQL/JRT.	123
16.3	Implied feature relationships of SQL/JRT.	123
Annex A	SQL Conformance Summary.	125
Annex B	Implementation-defined elements.	131
Annex C	Implementation-dependent elements.	135
Annex D	SQL feature taxonomy.	137
Annex E	Routines tutorial.	139
E.1	Technical components.	139
E.2	Overview.	140
E.3	Example Java methods: region and correctStates.	141
E.4	Installing region and correctStates in SQL.	141
E.5	Defining SQL names for region and correctStates.	143
E.6	A Java method with output parameters: bestTwoEmps.	144
E.7	A CREATE PROCEDURE best2 for bestTwoEmps.	146

E.8	Calling the best2 procedure.	146
E.9	A Java method returning a result set: orderedEmps.	147
E.10	A CREATE PROCEDURE rankedEmps for orderedEmps.	148
E.11	Calling the rankedEmps procedure.	149
E.12	Overloading Java method names and SQL names.	150
E.13	Java main methods.	152
E.14	Java method signatures in the CREATE statements.	152
E.15	Null argument values and the RETURNS NULL clause.	153
E.16	Static variables.	156
E.17	Dropping SQL names of Java methods.	157
E.18	Removing Java classes from SQL.	157
E.19	Replacing Java classes in SQL.	158
E.20	Visibility.	159
E.21	Exceptions.	159
E.22	Deployment descriptors.	160
E.23	Paths.	163
E.24	Privileges.	165
E.25	Information Schema.	166
Annex F	Types tutorial.	167
F.1	Overview.	167
F.2	Example Java classes.	167
F.3	Installing Address and Address2Line in an SQL system.	169
F.4	CREATE TYPE for Address and Address2Line.	170
F.5	Multiple SQL types for a single Java class.	171
F.6	Collapsing subclasses.	172
F.7	GRANT and REVOKE statements for data types.	174
F.8	Deployment descriptors for classes.	174
F.9	Using Java classes as data types.	175
F.10	SELECT, INSERT, and UPDATE.	176
F.11	Referencing Java fields and methods in SQL.	177
F.12	Extended visibility rules.	178
F.13	Logical representation of Java instances in SQL.	178
F.14	Static methods.	180
F.15	Static fields.	180
F.16	Instance-update methods.	181
F.17	Subtypes in SQL/JRT data.	183
F.18	References to fields and methods of null instances.	184
F.19	Ordering of SQL/JRT data.	185
Index.		189

TABLES

Page

1	System properties.	95
2	SQLSTATE class and subclass values.	122
3	Implied feature relationships of SQL/JRT.	123
4	Feature taxonomy for optional features.	137

INTRODUCTION

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

- 1) [Clause 1, “Scope”](#), specifies the scope of this part of ISO/IEC 9075.
- 2) [Clause 2, “Normative references”](#), identifies additional standards that, through reference in this part of ISO/IEC 9075, constitute provisions of this part of ISO/IEC 9075.
- 3) [Clause 3, “Definitions, notations, and conventions”](#), defines the notations and conventions used in this part of ISO/IEC 9075.
- 4) [Clause 4, “Concepts”](#), presents concepts used in the definition of Java routines and types.
- 5) [Clause 5, “Lexical elements”](#), defines a number of lexical elements used in the definition of Java routines and types.
- 6) [Clause 6, “Scalar expressions”](#), defines the elements of the language that produce scalar values.
- 7) [Clause 7, “Predicates”](#), defines the predicates of the language.
- 8) [Clause 8, “Additional common elements”](#), defines additional language elements that are used in various parts of the language.
- 9) [Clause 9, “Schema definition and manipulation”](#), defines the schema definition and manipulation statements associated with the definition of Java routines and types.
- 10) [Clause 10, “Access control”](#), defines facilities for controlling access to SQL-data.
- 11) [Clause 11, “Built-in procedures”](#), defines new built-in procedures used in the definition of Java routines and types.
- 12) [Clause 12, “Java topics”](#), defines the facilities supported by implementations of this part of ISO/IEC 9075 and the conventions used in deployment descriptor files.
- 13) [Clause 13, “Information Schema”](#), defines viewed tables that contain schema information.
- 14) [Clause 14, “Definition Schema”](#), defines base tables on which the viewed tables containing schema information depend.
- 15) [Clause 15, “Status codes”](#), defines SQLSTATE values related to Java routines and types.
- 16) [Clause 16, “Conformance”](#), defines the criteria for conformance to this part of ISO/IEC 9075.
- 17) [Annex A, “SQL Conformance Summary”](#), is an informative Annex. It summarizes the conformance requirements of the SQL language.
- 18) [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.
- 19) [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.

- 20) [Annex D, “SQL feature taxonomy”](#), is an informative Annex. It identifies features of the SQL language specified in this part of ISO/IEC 9075 by a numeric 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.
- 21) [Annex E, “Routines tutorial”](#), is an informative Annex. It provides a tutorial on using the features defined in this part of ISO/IEC 9075 for defining and using SQL-invoked routines based on Java static methods.
- 22) [Annex F, “Types tutorial”](#), is an informative Annex. It provides a tutorial on using the features defined in this part of ISO/IEC 9075 for defining and using SQL structured types based on Java classes.

In the text of this part of ISO/IEC 9075, Clauses begin a new odd-numbered page, and in [Clause 5, “Lexical elements”](#), through [Clause 16, “Conformance”](#), Subclauses begin a new page. Any resulting blank space is not significant.

X

NOTES

AUSTRALIAN STANDARD

Information technology — Database languages — SQL —

Part 13:

SQL Routines and Types Using the Java™ Programming Language (SQL/JRT)

1 Scope

This part of International Standard ISO/IEC 9075 specifies the 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 user-defined types. (Java is a registered trademark of Sun Microsystems, Inc.)

This page intentionally left blank.

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 JTC1 standards

[Framework] ISO/IEC 9075-1:2003, *Information technology — Database languages — SQL — Part 1: Framework (SQL/Framework)*.

[Foundation] ISO/IEC 9075-2:2003, *Information technology — Database languages — SQL — Part 2: Foundation (SQL/Foundation)*.

[OLB] ISO/IEC 9075-10:2003, *Information technology — Database languages — SQL — Part 10: Object Language Bindings (SQL/OLB)*.

[Schemata] ISO/IEC 9075-11:2003, *Information technology — Database languages — SQL — Part 11: Information and Definition Schemas (SQL/Schemata)*.

2.2 Other international standards

[Java] *The Java™ Language Specification, Second Edition*, Bill Joy (Editor), Guy Steele, James Gosling, and Gilad Bracha, Addison-Wesley, 2000, ISBN 0-201-31008-2.

[JVM] *The Java™ Virtual Machine Specification, Second Edition*, Tim Lindholm and Frank Yellin, Addison-Wesley, 1999, ISBN 0-201-43294-3.

[J2SE] *Java™ 2 Platform, Standard Edition, v1.3.1, API Specification*, <http://java.sun.com/j2se/1.3/docs/api/index.html>.

[Serialization] *Java™ Object Serialization Specification*, <http://java.sun.com/j2se/1.3/docs/guide/serialization/spec/serial-TOC.doc.html>.

[JavaBeans] *The JavaBeans™ 1.01 Specification*, <http://java.sun.com/products/javabeans/docs/spec.html>.

[JDBC] *JDBC™ 3.0 Specification, Final Release*, John Ellis & Linda Ho with Maydene Fisher, Sun Microsystems, Inc., October, 2001.

[RFC2368] RFC 2368, *The mailto URL schema*, P. Hoffman, L. Masinter, J. Zawinski, July, 1998. <http://www.ietf.org/rfc/rfc2368.txt>

[RFC2396] RFC 2396, *Uniform Resource Identifiers (URI): Generic Syntax*, T. Berners-Lee, R. Fielding, L. Maxinter, August, 1998. <http://www.ietf.org/rfc/rfc2396.txt>

3 Definitions, notations, and conventions

This Clause modifies Clause 3, “Definitions, notations, and conventions”, in ISO/IEC 9075-2.

3.1 Definitions

This Subclause modifies Subclause 3.1, “Definitions”, in ISO/IEC 9075-2.

3.1.1 Definitions taken from [Java]

For the purposes of this part of ISO/IEC 9075, the definitions of the following terms given in [Java] apply.

3.1.1.1 block

3.1.1.2 class declaration

3.1.1.3 class instance

3.1.1.4 class variable

3.1.1.5 field

3.1.1.6 instance initializer

3.1.1.7 instance variable

3.1.1.8 interface

3.1.1.9 local variable

3.1.1.10 nested class

3.1.1.11 package

3.1.1.12 static initializer

3.1.1.13 subpackage

3.1.2 Definitions taken from [JVM]

For the purposes of this part of ISO/IEC 9075, the definitions of the following terms given in [JVM] apply.

3.1.2.1 class file

3.1.2.2 Java Virtual Machine

3.1.3 Definitions provided in Part 13

For the purposes of this part of ISO/IEC 9075, in addition to those definitions adopted from other sources, the following definitions apply:

- 3.1.3.1 class file:** A file containing the compiled byte code for a Java class.
- 3.1.3.2 default connection:** A JDBC connection to the current SQL-implementation, SQL-session, and SQL-transaction established with the data source URL '`jdbc:default:connection`'. (See [RFC2368] and [RFC2396] for more details about URLs.)
- 3.1.3.3 deployment descriptor:** One or more SQL-statements that specify `<install actions>` and `<remove actions>` to be taken, respectively, by the `SQLJ.INSTALL_JAR` and `SQLJ.REMOVE_JAR` procedures and that are contained in a *deployment descriptor file*.
- 3.1.3.4 deployment descriptor file:** A text file containing deployment descriptors that is contained in a JAR, for which the JAR's manifest entry, as described by the `java.util.jar` section of [J2SE], specifies `SQLJDeploymentDescriptor: TRUE`.
- 3.1.3.5 external Java data type:** An SQL user-defined type defined with a `<user-defined type definition>` that specifies an `<external Java type clause>`.
- 3.1.3.6 external Java routine:** An external routine defined with an `<SQL-invoked routine>` that specifies `LANGUAGE JAVA` and either `PROCEDURE` or `FUNCTION`, or defined with a `<user-defined type definition>` that specifies an `<external Java type clause>`.
- 3.1.3.7 installed JAR:** A JAR whose existence has been registered with the SQL-environment and whose contents have been copied into that SQL-environment due to execution of one of the procedures `SQLJ.INSTALL_JAR` and `SQLJ.REPLACE_JAR`.
- 3.1.3.8 Java Archive (JAR):** A zip formatted file, as described by the `java.util.zip` section of [J2SE], containing zero or more Java `class` and `ser` files, and zero or more deployment descriptor files. JARs are a normal vehicle for distributing Java programs and the mechanism specified by this International Standard to provide the implementation of external Java routines and external Java data types to an SQL-environment.
- 3.1.3.9 JVM:** A Java Virtual Machine, as defined by [JVM].
- 3.1.3.10 ser file:** A file containing representations of Java objects in the form defined in [Serialization].
- 3.1.3.11 subject Java class:** The Java class uniquely identified by the combination of the class's *subject Java class name* and its containing JAR.
- 3.1.3.12 subject Java class name:** The fully-qualified package and class name of a Java class.
- 3.1.3.13 system class:** Any Java class provided by a conforming implementation of this part of ISO/IEC 9075 that can be referenced by an external Java routine or an external Java data type without that class having been included in an installed JAR.

3.2 Conventions

This Subclause modifies Subclause 3.3, “Conventions”, in ISO/IEC 9075-2.

3.2.1 Specification of built-in procedures

Built-in procedures are specified in terms of:

- **Function:** A short statement of the purpose of the procedure.
- **Signature:** A specification, in SQL, of the signature of the procedure. The only purpose of the signature is to specify the procedure name, parameter names, and parameter types. The manner in which these built-in procedures are defined is implementation-dependent.
- **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 invocation of the procedure. 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 the Signature, Access Rules, or General Rules can be referenced in other rules provided that they are defined before being referenced.

3.2.2 Specification of deployment descriptor files

Deployment descriptor files are specified in terms of:

- **Function:** A short statement of the purpose of the deployment descriptor file.
- **Model:** A brief description of the manner in which a deployment descriptor file is identified within its containing JAR.
- **Properties:** A BNF specification of the syntax of the contents of a deployment descriptor file.
- **Description:** A specification of the requirements and restrictions on the contents of a deployment descriptor file.
- **Conformance Rules:** A specification of how the element shall be supported for conformance to SQL.

This page intentionally left blank.

4 Concepts

This Clause modifies [Clause 4](#), “Concepts”, in ISO/IEC 9075-2.

4.1 The Java programming language

The Java programming language is a class-based, object-oriented language. This part of ISO/IEC 9075 uses the following Java concepts and terminology.

A *class* is the basic construct of Java programs, in that all executable Java code is contained in a Java class definition. A class is declared by a *class declaration* that specifies a possibly empty set consisting of zero or more fields, zero or more methods, zero or more nested classes, zero or more interfaces, zero or more instance initializers, zero or more static initializers, and zero or more constructors.

The scope of a variable is a class, an instance of the class, or a method of the class. The scope of a variable that is declared *static* is the class, and the variable is called a *class variable*. The scope of each other variable declared in the class is instances of the class, and such a variable is called an *instance variable*. Class variables and instance variables of a class are called *fields* of that class. The scope of a variable declared in a method is the *block* or Java **for** statement in which it is declared in that method, and the variable is called a *local variable*.

A *class instance* consists of an instance of each instance variable declared in the class and each instance variable declared in the superclasses of the class. Class instances are strongly typed by the class name. The operations available on a class instance are those defined for its class.

With the exception of the class `java.lang.Object`, each class is declared to *extend* (at most) one other class; a class not explicitly declared to extend another class implicitly extends `java.lang.Object`. The declared class is a *direct subclass* of the class that it extends; the class that it extends is the *direct superclass* of the declared class.

Class *B* is a *subclass* of class *A* if *B* is a direct subclass of *A*, or if there exists some class *C* such that *B* is a direct subclass of *C* and *C* is a subclass of *A*. Likewise, class *B* is a *superclass* of class *A* if *B* is a direct superclass of *A*, or if there exists some class *C* such that *B* is a direct superclass of *C* and *C* is a superclass of *A*. A subclass has all of the fields and methods of its superclasses and an instance of class *B* may be used wherever an instance of a superclass of *B* is permitted.

A *method* is an executable routine. A method can be declared *static*, in which case it is called a *class method*; otherwise, it is called an *instance method*. A class method can be referenced by qualifying the method name with either the class name or the name of an instance of the class. An instance method is referenced by qualifying the method name with a Java expression that results in an instance of the class or, in the case of a constructor, with “new”. The method body of an instance method can reference its class variables, instance variables, and local variables.

The *Java method signature* of a method consists of the method name and the number of parameters and their data types.

A *package* consists of zero or more classes, zero or more interfaces, and zero or more *subpackages* (a subpackage is a package within a package); each package provides its own name space and classes within a package are able to refer to other classes in the same package, including classes not referenceable from outside the package. Every class belongs to exactly one package, either an explicitly specified named package or the anonymous default package. A class can specify Java `import` statements to refer to other named packages whose classes can then be referenced within the class without package qualification.

A class, field, or methods can be declared as *public*, *private*, or *protected*. A public variable or method can be accessed by any method. A private variable or method can only be referenced by methods in the same class. A protected variable or method can only be referenced by methods of the same class or subclasses thereof. A method that is not declared as public, private, or protected can only be called by methods declared by classes in the same package.

An *interface* is a Java construct consisting of a set of method signatures. An interface can be implemented by zero or more classes, a class can be declared to implement zero or more interfaces, and a class is required to have methods with the signatures specified by all of its declared interfaces.

The Java *Serializable* interface, `java.io.Serializable`, as described in [J2SE], defines a transformation between a Java instance and a `java.io.OutputStream` or `java.io.InputStream`, as defined by the `java.io.OutputStream` and `java.io.InputStream` sections of [J2SE] respectively, writing a persistent representation of an instance of a Java object and reading a persistent representation of an instance of a Java object. This transformation retains sufficient information to identify the most specific class of the instance and to reconstruct the instance.

The Java *SQLData* interface, `java.sql.SQLData`, as described in [JDBC] and [J2SE], defines a transformation between a Java instance and an SQL user-defined data type.

The source for a Java class is normally stored in a *Java file* with the file-type “java”, e.g., `myclass.java`. Java is normally compiled to a byte coded instruction set that is portable to any platform supporting Java. A file containing such byte code is normally stored in a *class file* with the file-type “class”, e.g., `myclass.class`.

A set of class files can be assembled into a *Java archive* file, or *JAR* (usually with a file extension of “.jar”). A JAR is a zip formatted file containing a set of Java class files. JARs are the normal vehicle for distributing Java programs.

4.2 SQL-invoked routines

This Subclause modifies Subclause 4.27, “SQL-invoked routines”, in ISO/IEC 9075-2.

4.2.1 Overview of SQL-invoked routines

This Subclause modifies Subclause 4.27.1, “Overview of SQL-invoked routines”, in ISO/IEC 9075-2.

Replace the lead text of the 9th paragraph A static SQL-invoked method, whether or not it is an external Java routine, satisfies the following conditions:

4.2.2 Characteristics of SQL-invoked routines

This Subclause modifies Subclause 4.27.2, “Characteristics of SQL-invoked routines”, in ISO/IEC 9075-2.

Insert after 2nd paragraph External routines appear in two seemingly similar, but fundamentally differing, forms, where the key differentiator is whether or not the external routine's routine descriptor specifies that the body of the SQL-invoked routine is written in Java. When the body of the SQL-invoked routine is written in Java, the external routine is an *external Java routine*; some differences from other external routines include:

- For any other external routine, the *executable form* (such as a dynamic link library or some run-time interpreted form) of that routine exists externally to the SQL-environment's catalogs; for an external Java routine, the executable form is provided by a specified subject Java routine that exists in the SQL-environment's catalogs in an installed JAR.
- Because an installed JAR is not required to be completely self-contained (*i.e.*, to have no references to Java classes outside of itself), a mechanism is provided to allow a subject Java class to reference classes defined by class files contained in its installed JAR or in other installed JARs. See Subclause 8.2, “<SQL Java path>”.

NOTE 1 — Once an external Java routine has been created, its use in SQL statements executed by the containing SQL-environment is similar to that of other external routines.

Replace the 4th paragraph SQL-invoked routines are invoked differently depending on their form. SQL-invoked procedures are invoked by <call statement>s. SQL-invoked regular functions are invoked by <routine invocation>s. Instance SQL-invoked methods are invoked by <method invocation>s, while SQL-invoked constructor methods are invoked by <new invocation>s and static SQL-invoked methods are invoked by <static method invocation>s. An invocation of an SQL-invoked routine specifies the <routine name> of the SQL-invoked routine and supplies a sequence of argument values corresponding to the <SQL parameter declaration>s of the SQL-invoked routine. A *subject routine* of an invocation is an SQL-invoked routine that may be invoked by a <routine invocation>. After the selection of the subject routine of a <routine invocation>, the SQL arguments are evaluated and the SQL-invoked routine that will be executed is selected. If the subject routine is an instance SQL-invoked method that is not an external Java routine, then the SQL-invoked routine that is executed is selected from the set of overriding methods of the subject routine. (The term “set of overriding methods” is defined in the General Rules of Subclause 10.4, “<routine invocation>”, in ISO/IEC 9075-2.) The overriding method that is selected is the overriding method with a subject parameter the type designator of whose declared type precedes that of the declared type of the subject parameter of every other overriding method in the type precedence list of the most specific type of the value of the SQL argument that corresponds to the subject parameter. (See the General Rules of Subclause 10.4, “<routine invocation>”, in ISO/IEC 9075-2.) If the instance SQL-invoked method is an external Java routine, the term “set of overriding methods” is not applicable; for such methods, the capabilities provided by overriding methods duplicate Java's own mechanisms and the subject routine executed is the one that would be invoked when no overriding methods are specified. If the subject routine is not an SQL-invoked method, then the SQL-invoked routine executed is that subject routine. After the selection of the SQL-invoked routine for execution, the values of the SQL arguments are assigned to the corresponding SQL parameters of the SQL-invoked routine and its <routine body> is executed. If the SQL-invoked routine is an SQL routine, then the <routine body> is an <SQL procedure statement> that is executed according to the General Rules of Subclause 13.5, “<SQL procedure statement>”, in ISO/IEC 9075-2. If the SQL-invoked routine is an external routine, then the <routine body> identifies a program written in some standard programming language other than SQL that is executed according to the rules of that programming language.

Replace the 6th paragraph If the SQL-invoked routine is an external routine, then an effective SQL parameter list is constructed before the execution of the <routine body>. The effective SQL parameter list has different entries depending on the parameter passing style of the SQL-invoked routine. The value of each entry in the effective SQL parameter list is set according to the General Rules of [Subclause 8.3](#), “<routine invocation>”. When the SQL-invoked routine is not an external Java routine, the values in the effective SQL parameter list are passed to the program identified by the <routine body> according to the rules of [Subclause 13.6](#), “Data type correspondences”, in ISO/IEC 9075-2; when the SQL-invoked routine is an external Java routine, values in the effective SQL parameter list are passed to the program identified by <routine body> according to the rules of [Subclause 4.5](#), “Parameter mapping”. After the execution of that program, if the parameter passing style of the SQL-invoked routine is SQL, then the SQL-implementation obtains the values for output parameters (if any), the value (if any) returned from the program, the value of the exception data item, and the value of the message text (if any) from the values assigned by the program to the effective SQL parameter list. If the parameter passing style of the SQL-invoked routine is JAVA, then such values are obtained from the values assigned by the program to the effective SQL parameter list and the uncaught Java exception (if any). If the parameter passing style of the SQL-invoked routine is GENERAL, then such values are obtained in an implementation-defined manner.

4.3 Java class name resolution

Typical JVMs provide a *class name resolution*, or search path, mechanism based on an environmental variable called CLASSPATH. When a JVM encounters a previously unseen reference to a class, the members of the list of directories and JARs appearing in the classpath are examined in order until either the class is found or the end of the list is reached. Failure to locate a referenced class is a runtime error that will often cause the application that experiences it to terminate.

When a JVM is transitioned to being effectively within an SQL environment, the problem of managing the JVM's class name resolution continues to exist, but with a change in emphasis. To allow the creators of Java applications a greater degree of control over class name resolution, and the added security associated with that control, a classpath-like mechanism is defined to be a property of the JARs containing the Java applications, rather than as an environmental variable of the current session (such as, for example, CURRENT_PATH for dynamic statements). Therefore, if, while an external Java routine is being executed, a previously unseen class reference is encountered, that class is first searched for in the JAR containing the definition of the currently executing class, and, if it is not found, the class will be sought in the manner specified by the SQL-Java path associated with that JAR (if any).

An SQL-Java path specifies how a JVM resolves a class name when a class within a JAR references a class that is not a system class or not in the same JAR. `SQLJ.ALTER_JAVA_PATH` is used to associated an SQL-Java path with a JAR. An SQL-Java path is a list of (referenced item, referenced JAR) pairs. A referenced item can be either a class, a package, or '*' to specify the entire JAR. The SQL-Java path list is searched in the order the pairs are specified. For each (referenced item, referenced JAR) pair (*RI*, *RJ*):

- If *RI* is the class name, then the class shall be defined in *RJ*. If it is not, an exception condition is raised.
- If *RI* is the package of the class being resolved, then the class shall be defined in *RJ*. If it is not, an exception condition is raised.
- If *RI* is '*' and the class is defined in *RJ*, then that resolution is used; otherwise, subsequent pairs are tested.

4.4 SQL result sets

Cursors, or SQL result sets, appear to Java applications in two forms; the first, as an object of a class that implements the interface `java.sql.ResultSet` as defined in [JDBC] and [J2SE], referred to as a *JDBC ResultSet*; the second, as an object of a class that implements the interface `sqlj.runtime.ResultSetIterator` as defined by ISO/IEC 9075-10, referred to as an *SQLJ Iterator*.

In ISO/IEC 9075-2, SQL-invoked procedures are declared to be able to return zero or more dynamic result sets, referred to as *result set cursors*. To be a returned result set cursor, a cursor's declaration shall specify `WITH RETURN`, and the cursor shall be open at the point that the SQL-invoked procedure exits. While external Java routines that are SQL-invoked procedures can likewise be declared to return zero or more dynamic result sets, in some other respects, this part of ISO/IEC 9075's treatment of result set cursors differs from that of ISO/IEC 9075-2.

In a Java application, all JDBC ResultSets and SQLJ Iterators are implicitly result set cursors, that is, their underlying cursor declarations implicitly specify `WITH RETURN`. So, in this part of ISO/IEC 9075, to actually be a returned result set cursor it is not sufficient that the corresponding JDBC ResultSet's or SQLJ Iterator's underlying cursor be open when the SQL-invoked procedure exits; the JDBC ResultSet or SQLJ Iterator shall also have been explicitly assigned to a parameter of the subject Java routine that represents an output parameter. As discussed in Subclause 4.5, “Parameter mapping”, and Subclause 8.3, “<routine invocation>”, output parameters are represented to a subject Java routine as the first element of a one dimensional array of a Java data type that can be mapped to an SQL data type. For dynamic result sets, the array shall be of a class that implements the interface `java.sql.ResultSet` or the interface `sqlj.runtime.ResultSetIterator`, the JDBC ResultSet or SQLJ Iterator shall have been explicitly assigned to the first element of that array, and that JDBC ResultSet or SQLJ Iterator shall not have been closed.

It is important to note that this difference in how a result set cursor becomes a returned result set cursor is invisible to the calling application. As described in Subclause 8.3, “<routine invocation>”, the calling application will be returned zero or more dynamic result sets in the order in which the cursors were opened, just as in ISO/IEC 9075-2; the order of the parameters in the subject Java routine does not impact the order in which the calling application accesses the returned result sets.

4.5 Parameter mapping

Let *ST* be some SQL data type and let *JT* be some Java data type.

ST and *JT* are *simply mappable* if and only if *ST* and *JT* are specified respectively in the first and second columns of some row of the *Data type conversion tables*, Table B.1, entitled “JDBC Types mapped to Java Types”, in [JDBC]. The Java data type *JT* is the *corresponding Java data type* of *ST*.

ST and *JT* are *object mappable* if and only if *ST* and *JT* are specified respectively in the first and second columns of some row of the *Data type conversion tables*, Table B.3, entitled “Mapping from JDBC Types to Java ObjectTypes”, in [JDBC], or if the descriptor of *ST* specifies that it is an external Java data type and the descriptor specifies *JT* as the <Java class name> in the <jar and class name>.

ST and *JT* are *output mappable* if and only if:

- *JT* is a one dimensional array type with an element data type *JT2* (that is, *JT* is “*JT2[]*”) and *ST* is either simply mappable to *JT2* or object mappable to *JT2*.
- *JT* is `java.lang.StringBuffer` and its corresponding parameter in the augmented SQL parameter declaration list is the save area data item.

An SQL array type with an element data type *ST* and *JT* are *array mappable* if and only if *JT* is a one dimensional array type with an element data type *JT2* and *ST* is either simply mappable to *JT2* or object mappable to *JT2*.

ST and *JT* are *mappable* if and only if *ST* and *JT* are simply mappable, object mappable, output mappable, or array mappable.

A Java data type is *mappable* if and only if it is mappable to some SQL data type.

A Java class is *result set oriented* if and only if it is either:

- A class that implements the Java interface `java.sql.ResultSet`.
- A class that implements the Java interface `sqlj.runtime.ResultSetIterator`.

NOTE 2 — These classes are generated by iterator declarations (`#sql iterator`) as specified in ISO/IEC 9075-10.

A Java data type is *result set mappable* if and only if it is a one-dimensional array type with an element type that is a result set oriented class.

A Java method with *M* parameters is *mappable* (to SQL) if and only if, for some *N*, $0 \text{ (zero)} \leq N \leq M$, the data types of the first *N* parameters are mappable, the last *M–N* parameters are result set mappable, and the result type is either simply mappable, object mappable, or `void`.

A Java method is *visible* in SQL if and only if it is public and mappable. In addition, to be visible, a Java method shall be static if used as the external Java routine of an SQL-invoked procedure or an SQL-invoked regular function.

A Java class is *visible* in SQL if and only if it is public and mappable.

[JDBC] contains JDBC's SQL to Java data type mappings defined in the JDBC type mapping tables. If *ST* is an external Java data type that appears in the INFORMATION_SCHEMA.USER_DEFINED_TYPES view, then let *JT* be *ST*'s descriptor's <Java class name> in its <jar and class name>. JDBC's data type mapping tables are effectively extended. A row (*ST*, *JT*) is considered to be an additional row in Table B.3, *Mapping from JDBC Types to Java Object Types*, and a row (*JT*, *ST*) is considered to be an additional row in Table B.4, *Mapping from Java Object Types to JDBC Types*.

4.6 Unhandled Java exceptions

Java exceptions that are thrown during execution of a Java method in SQL can be caught, or handled, within Java; if this is done, then those exceptions do not affect SQL processing. All Java exceptions that are uncaught when a Java method called from SQL completes appear in the SQL-environment as SQL exception conditions.

The message text may be specified in the Java exception specified in the Java `throw` statement. If the Java exception is an instance of `java.sql.SQLException`, or a subtype of that type, then it may also specify an SQLSTATE value. If that exception specifies an SQLSTATE value, the first two characters of that SQLSTATE shall be '38'. If that exception does not specify an SQLSTATE value, then the default SQL exception condition

for an uncaught Java exception is raised. See Subclause 15.1, “Class and subclass values for uncaught Java exceptions”.

When a Java method executes an SQL statement, any exception condition raised in the SQL statement will be raised in the Java method as a Java exception that is specifically the `java.sql.SQLException` subclass of the Java class `java.lang.Exception`. For portability, a Java method called from SQL, that itself executes an SQL statement and that catches an `SQLException` from that inner SQL statement, should re-throw that `SQLException`.

4.7 Data types

This Subclause modifies Subclause 4.1, “Data types”, in ISO/IEC 9075-2.

4.7.1 Non-predefined and non-SQL types

This Subclause modifies Subclause 4.1.3, “Non-predefined and non-SQL types”, in ISO/IEC 9075-2.

Replace the 6th paragraph Each host language has its own data types, which are separate and distinct from SQL data types, even though similar names may be used to describe the data types. Mappings of SQL data types to data types in host languages are described in Subclause 11.50, “<SQL-invoked routine>”, in ISO/IEC 9075-2, in Subclause 20.1, “<embedded SQL host program>”, in ISO/IEC 9075-2, and in Subclause 8.1, “<embedded SQL host program>”, in ISO/IEC 9075-10. Not every SQL data type has a corresponding data type in every host language.

4.8 User-defined types

This Subclause modifies Subclause 4.7, “User-defined types”, in ISO/IEC 9075-2.

4.8.1 Introduction to user-defined types

This Subclause modifies Subclause 4.7.1, “Introduction to user-defined types”, in ISO/IEC 9075-2.

Insert after 1st paragraph User-defined types appear in two seemingly similar, but fundamentally differing, forms in which the key differentiator is whether or not the create type statement for the user-defined type specifies an external language of “JAVA”. When an external language of JAVA is specified, the user-defined type is an *external Java data type* and the create type statement defines a mapping of the user-defined type's attributes and methods directly to the public attributes and methods of a *subject Java class*. This is different from user-defined types that are not external Java data types. The differences include:

- For every other user-defined type, there is no requirement for an association with an underlying class; each method of a user-defined type that is not an external Java data type can be written in a different language

(for example, one method could be written in SQL and another written in Fortran). Such user-defined types cannot have methods written in Java. By contrast, all methods of an external Java data type shall be written in Java, (implicitly) have a parameter style of JAVA, and be defined in the associated Java class or one of its superclasses.

- For every other user-defined type, there is no explicit association between a user-defined type's attributes and any external representation of their content. In addition, the mapping between a user-defined type's methods and external methods is made over time by subsequent CREATE METHOD statements. By contrast, for external Java data types, the association between the user-defined type's attributes and methods and the public attributes and methods of a subject Java class is specified by the create type statement.
- For external Java data types, the mechanism used to convert the SQL-environment's representation of an instance of a user-defined type into an instance of a Java class is specified in the USING <interface specification> clause. Such conversions are performed, for example, when an external Java data type is specified as a (subject) parameter in a method or function invocation, or when a Java object returned from a method or function invocation is stored in a column declared to be an external Java data type. <interface specification> can be either SERIALIZABLE, specifying the Java-defined interface `java.io.Serializable` (not to be confused with the isolation level of SERIALIZABLE), or SQLDATA, specifying the JDBC-defined interface `java.sql.SQLData`. See Subclause 9.4, “<user-defined type definition>”.
- For every other user-defined type, there is no explicit support of static attributes. For external Java data types, the <user-defined type definition> is allowed to include <static field method spec>s that define observer methods against specified static attributes of the subject Java class.

The scope and persistence of any modifications to static attributes made during the execution of a Java method is implementation-dependent.

- For every other user-defined type, the implementation of every method that isn't an SQL routine exists externally to the SQL-environment. For external Java data types, the implementation of the methods is provided by a specified subject Java class that exists within the SQL-environment in an *installed JAR*.
- External Java data types may only be structured types, not distinct types.
- Support for the specification of overriding methods is not provided for methods that are external Java routines.

NOTE 3 — Once an external Java data type has been created, its use in SQL statements executed by the containing SQL-implementation is similar to that of other user-defined types.

4.8.2 User-defined type descriptor

This Subclause modifies Subclause 4.7.2, “User-defined type descriptor”, in ISO/IEC 9075-2.

Augment the 4th list element in the 1st paragraph by inserting “COMPARABLE,”

- Insert into the list in the 1st paragraph An indication of whether the user-defined type is an external Java data type.

Insert following the 1st paragraph If the user-defined type is an external Java data type, then the user-defined type descriptor also includes:

- The <jar and class name> of the user-defined type.

- The <interface specification> of `SERIALIZABLE` or `SQLDATA`.
- The attribute descriptor of every originally-defined attribute and every inherited attribute of the user-defined type.
- If <method specification list> is specified, then, for each <method specification> contained in <method specification list>, a *method spec descriptor* that includes:
 - The <method name>.
 - The <specific method name>.
 - The <SQL parameter declaration list>.
 - The <returns data type>, and indication of `SELF AS RESULT`.
 - The <result cast from type>, if any.
 - The package, class, and name of the Java routine corresponding to this method and, if specified, its signature.
 - An indication of whether `STATIC` or `CONSTRUCTOR` is specified.
 - If `STATIC` is specified, then an indication of whether this is a static field method.
 - If this is a static field method, then the <Java field name> of the static field and the <Java class name> of the class that declares that static field.
 - An indication of whether the method is deterministic.
 - An indication of whether the method possibly writes SQL data, possibly reads SQL data, possibly contains SQL, or does not possibly contain SQL.
 - An indication of whether the method should not be invoked if any argument is the null value, in which case the value of the method is the null value.

If the user-defined type is not an external Java data type, then the user-defined type descriptor also includes:

- An indication of whether the user-defined type is a structured type or a distinct type.
- If the representation is a predefined data type, then the descriptor of that type; otherwise, the attribute descriptor of every originally-defined attribute and every inherited attribute of the user-defined type.
- If the <method specification list> is specified, then, for each <method specification> contained in <method specification list>, a *method specification descriptor* that includes:
 - The <method name>.
 - The <specific method name>.
 - The <SQL parameter declaration list> augmented to include the implicit first parameter with parameter name `SELF`.
 - The <language name>.
 - If the <language name> is not `SQL`, then the <parameter style>.
 - The <returns data type>.

- The <result cast from type>, if any.
- An indication as to whether the <method specification> is an <original method specification> or an <overriding method specification>.
- If the <method specification> is an <original method specification>, then an indication of whether STATIC or CONSTRUCTOR is specified.
- An indication whether the method is deterministic.
- An indication whether the method possibly writes SQL data, possibly reads SQL data, possibly contains SQL, or does not possibly contain SQL.
- An indication whether the method should not be invoked if any argument is the null value, in which case the value of the method is the null value.

NOTE 4 — The characteristics of an <overriding method specification> other than the <method name>, <SQL parameter declaration list>, and <returns data type> are the same as the characteristics for the corresponding <original method specification>.

4.8.3 User-defined type comparison and assignment

This Subclause modifies Subclause 4.7.6, “User-defined type comparison and assignment”, in ISO/IEC 9075-2.

Replace 5th paragraph Let *comparison function* of a user-defined type T_a be the ordering function included in the user-defined type descriptor of the comparison type of T_a , if any.

Replace 6th paragraph Two values $V1$ and $V2$ whose most specific types are user-defined types $T1$ and $T2$ are comparable if and only if $T1$ and $T2$ are in the same subtype family and each have some comparison type $CT1$ and $CT2$, respectively. $CT1$ and $CT2$ constrain the comparison forms and comparison categories of $T1$ and $T2$ to be the same and to be the same as those of all their supertypes. If the comparison category is COMPARABLE, then no comparison functions shall be specified for $T1$ and $T2$. If the comparison category is either STATE or RELATIVE, then the comparison functions of $T1$ and $T2$ are constrained to be equivalent. If the comparison category is MAP, they are not constrained to be equivalent.

4.8.4 Accessing static fields

The fields of a Java class can be defined to be either *static* or *non-static*. Static fields of a Java class can additionally be specified to be *final*, which makes them read-only. In Java, non-final fields are allowed to be updated.

SQL's <user-defined type definition> does not include a facility for specifying attributes to be STATIC. This is, in part, because of the difficulty in specifying the scope, persistence, and transactional properties of static attributes in a database environment. An external Java data type's <user-defined type definition> does, however, provide a mechanism for read-only access to the values of Java static fields. The <static field method spec> clause defines a method name for a method with no parameters; its <external variable name clause> specifies the name of a static field of the subject Java class or a superclass of the subject Java class. A static field method is invoked in the normal manner for STATIC methods and returns the value of the specified Java static field. Whether final or non-final, SQL provides no mechanism for updating the values of Java static fields.

4.8.5 Converting objects between SQL and Java

While application programmers or end users manipulating Java objects in the database through SQL statements need not be aware of the specific mechanism used to achieve that conversion, the developer of the Java class itself needs to prepare for it in the form of implementing special Java interfaces (*i.e.*, `java.io.Serializable` or `java.sql.SQLData`). <user-defined type definition> introduces a clause for specifying the interface for converting object state information between the SQL database and Java in the scope of SQL statements. As mentioned above, a conversion from SQL to Java can potentially take place when an object that has been persistently stored in the SQL database is accessed from inside an SQL statement to retrieve attribute (or field) values, or to invoke a method on the object, or when the object is used as an input argument in the invocation of a method. A conversion in the opposite direction, from Java to SQL, may be required when a newly created or modified object, or an object that is the return value of a method invocation, needs to be persistently stored in the database.

This International Standard supports these options to specify object state conversion in the <external Java type clause>:

- If the <user-defined type definition> specifies an <interface specification> of `SERIALIZABLE`, then the Java interface `java.io.Serializable` is used for object state conversion.
- If the <user-defined type definition> specifies an <interface specification> of `SQLDATA`, then the Java interface `java.sql.SQLData` defined in [JDBC] and [J2SE] is used for object state conversion.
- If the <user-defined type definition> does not specify an <interface specification>, then it is implementation-defined whether the Java interface `java.io.Serializable` or the Java interface `java.sql.SQLData` will be used for object state conversion.

4.8.5.1 SERIALIZABLE

If the <interface specification> of a <user-defined type definition> specifies `SERIALIZABLE`, then object state communication is based on the Java interface `java.io.Serializable`. The Java class referenced in the <external Java class clause> of the <user-defined type definition> shall specify “implements `java.io.Serializable`” and shall provide a no-arg constructor.

In this case, the SQL object state that is stored persistently and made available to methods of the SQL type is defined entirely by the Java serialized object state. The attributes defined for the SQL type shall correspond to public fields of the corresponding Java class, which shall be listed in the <external Java attribute clause> of each attribute. Consequently, the SQL attributes define access to those portions of the object state that are intended to become visible inside SQL statements, but might not comprise the complete state of the object (which may include additional fields in the Java class).

4.8.5.2 SQLDATA

If the <interface specification> of a <user-defined type definition> specifies `SQLDATA`, then object state communication is based on the Java interface `java.sql.SQLData` defined in [JDBC] and [J2SE]. The Java

class referenced in the <external Java class clause> of the <user-defined type definition> shall specify “implements java.sql.SQLData” and shall provide a no-arg constructor.

In this case, only the attributes defined in the statement comprise the complete state of the SQL object type. Additional public or private attributes defined in the Java class do not become part of the object state defined by this part of ISO/IEC 9075. The Java object representation may be entirely different from the SQL object attributes, if desired. For example, an SQL Point type may define a geometric point in terms of cartesian coordinates, while the corresponding Java class defines it using polar coordinates. The only requirement to be met by the implementor of the Java class is that the implementations of the java.sql.SQLData methods readSQL and writeSQL read and write the attributes in the same order in which they are defined in the <user-defined type definition>.

To improve portability, it is possible to also specify <external Java attribute clause>s for SQL attributes, even if an <interface specification> of SQLData is specified. However, the <external Java attribute clause>s are ignored in this case, because they are not needed for implementing attribute access in SQL or for converting objects between SQL and Java.

4.8.5.3 Developing for portability

The following guidelines provide maximum portability of Java classes across different implementations of this part of ISO/IEC 9075 that may not support both the SERIALIZABLE and the SQLData options:

- The Java class used for implementing the SQL type should implement both java.io.Serializable and java.sql.SQLData.
- The Java class should define the complete object state that needs to become persistent or has to be preserved across invocations as public Java fields.
- The EXTERNAL NAMES of the SQL attributes should be specified.

The <interface using clause> should be omitted in the <user-defined type definition>, so that an implementation that does not support both interfaces can default to the interface that it supports.

4.9 Built-in procedures

This part of ISO/IEC 9075 differs slightly from other parts of ISO/IEC 9075 in its treatment of the schema object introduced to install the external Java routines and external Java data types in an SQL-environment — that is, in its treatment of JARs. Rather than define new SQL-schema statements that (for example) add or drop JARs using optional clauses to cause execution of their contained deployment descriptors, this International Standard introduces a set of four built-in procedures and a new schema in which those procedures are defined.

The new schema — named SQLJ — is, like the schema named INFORMATION_SCHEMA, defined to exist in all catalogs of an SQL system that implements this part of ISO/IEC 9075, and to contain all of the built-in procedures defined in this part of ISO/IEC 9075.

Built-in procedures defined in this part of ISO/IEC 9075 are:

- SQLJ.INSTALL_JAR — to load a set of Java classes in an SQL system.

- `SQLJ.REPLACE_JAR` — to supersede a set of Java classes in an SQL system.
- `SQLJ.REMOVE_JAR` — to delete a previously installed set of Java classes.
- `SQLJ.ALTER_JAVA_PATH` — to specify a path for name resolution within Java classes.

4.10 Privileges

This Subclause modifies *Subclause 4.34.2, “Privileges”*, in ISO/IEC 9075-2.

Augment the list in the 1st paragraph by inserting “JAR,”

Replace 2nd paragraph, 1st bullet

- The identification of the base table, view, column, domain, character set, collation, transliteration, user-defined type, table/method pair, trigger, SQL-invoked routine module, sequence generator, or JAR that the descriptor describes.

Replace 8th paragraph A privilege descriptor with an <action> of USAGE is called a *usage privilege descriptor* and identifies the existence of a privilege on the domain, user-defined type, character set, collation, transliteration, sequence generator, or JAR identified by the privilege descriptor.

Insert after 8th paragraph The privileges for facilities defined in this part of ISO/IEC 9075 are as follows:

- The privileges required to invoke the `SQLJ.INSTALL_JAR`, `SQLJ.REPLACE_JAR`, and `SQLJ.REMOVE_JAR` procedures are implementation-defined.
NOTE 5 — This is similar to the implementation-defined privileges required for creating a schema.
- Only the owner of the JAR is permitted to invoke the `SQLJ.ALTER_JAVA_PATH` procedure and the owner shall also have the USAGE privilege on each JAR referenced in the path argument.
- Invocations of <SQL-invoked routine> and <drop routine statement> to define and drop external Java routines are governed by the normal Access Rules for SQL-schema statements.
- Invocations of Java methods referenced by SQL names are governed by the normal EXECUTE privilege on SQL routine names.

It is implementation-defined whether a Java method called by an SQL name executes with “definer's rights” or “invoker's rights” — that is, whether it executes with the user-name of the user who performed the <SQL-invoked routine> or the user-name of the current user.

4.11 JARs

A JAR is a zip-formatted file containing a set of Java `class` and `ser` files and optionally a deployment descriptor file. Installed JARs provide the implementation of external Java routines and external Java data types to an SQL-environment.

JARs are created outside the SQL-environment. They are copied into the SQL-environment by the `SQLJ.INSTALL_JAR` procedure. No subsequent SQL statement or procedure modifies an installed JAR in

any way, other than to remove it from the SQL-environment, to replace it in its entirety, or to alter its SQL-Java path. In particular, no SQL operation adds classes to a JAR, removes classes from a JAR, or replaces classes in a JAR. The reason for this “no modification” principle for installed JAR is that JARs are often signed, and often contain *manifest* data that might be invalidated by modification of JARs by the SQL-environment.

Each installed JAR is represented by a *JAR descriptor*. A JAR descriptor contains:

- The catalog name, schema name, and JAR identifier of the JAR.
- The SQL-Java path of the JAR.

4.11.1 Deployment descriptor files

When a JAR is installed, one or more <SQL-invoked routine>s that define external Java routines shall be executed before the static methods of its contained Java classes can be used as SQL-invoked routines, and one or more <user-defined type definition>s shall be executed before its contained classes can be used as user-defined types. In addition, <grant privilege statement>s may be required to define privileges for newly created SQL-invoked routines and user-defined types. Later, when a JAR is removed, corresponding <drop routine statement>s, <drop data type statement>s, and <revoke statement>s shall be executed.

If a JAR is to be installed in several SQL implementations, the <SQL-invoked routine>s, <user-defined type definition>s, <user-defined ordering definition>s, <grant privilege statement>s, <drop routine statement>s, <drop data type statement>s, <drop user-defined ordering statement>s, and <revoke statement>s will often be the same for each implementation. To assist the automation of repeated installations, deployment descriptor files contain the variants of SQL-schema statements defined in this part of ISO/IEC 9075. These statements are grouped into multi-statement *install actions* and *remove actions* respectively executed by SQLJ . INSTALL_JAR and SQLJ . REMOVE_JAR procedures when deployment is requested. In addition, an implementation-defined *implementor block* is provided to allow specification of custom install and remove actions. Since the SQL-schema statements refer to their containing JAR in the <SQL-invoked routine>s and <user-defined type definition>s, within a deployment descriptor file, the JAR name “thisjar” is used as a place holder JAR name for the containing JAR.

This part of ISO/IEC 9075 provides a new mechanism to execute its variants of SQL-schema statements, namely by requesting deployment during invocation of SQLJ . INSTALL_JAR and SQLJ . REMOVE_JAR procedures. A conforming SQL-implementation is required to support either deployment descriptor based execution of its SQL-schema statements (Feature J531, “Deployment”) or another standard statement execution mechanism such as direct invocation or embedded SQL (Feature J511, “Commands”); a conforming SQL-implementation is not required to support both mechanisms.

5 Lexical elements

This Clause modifies Clause 5, “Lexical elements”, in ISO/IEC 9075-2.

5.1 <token> and <separator>

This Subclause modifies Subclause 5.2, “<token> and <separator>”, in ISO/IEC 9075-2.

Function

Specify lexical units (tokens and separators) that participate in SQL language.

Format

```
<non-reserved word> ::=
    !! All alternatives from ISO/IEC 9075-2
    | COMPARABLE

    | INTERFACE

    | JAVA

    | SQLDATA

<reserved word> ::=
    !! All alternatives from ISO/IEC 9075-2
    | JAR
```

Syntax Rules

No additional Syntax Rules.

Access Rules

No additional Access Rules.

General Rules

No additional General Rules.

Conformance Rules

No additional Conformance Rules.

5.2 Names and identifiers

This Subclause modifies *Subclause 5.4, “Names and identifiers”*, in ISO/IEC 9075-2.

Function

Specify names.

Format

```
<jar name> ::= [ <schema name> <period> ] <jar id>
<jar id> ::= <identifier>
<Java class name> ::= [ <packages> <period> ] <class identifier>
<jar and class name> ::= <jar id> <colon> <Java class name>
<qualified Java field name> ::= [ <Java class name> <period> ] <Java field name>
<packages> ::= <package identifier> [ <period> <package identifier> ]...
<package identifier> ::= <Java identifier>
<class identifier> ::= <Java identifier>
<Java field name> ::= <Java identifier>
<Java method name> ::= <Java identifier>
<Java identifier> ::= !! See the Syntax Rules
```

Syntax Rules

- 1) Insert this SR <Java identifier> shall be a valid identifier according to the rules of Java parsing and lexical analysis.
NOTE 6 — The rules of Java parsing and lexical analysis are specified in [Java].
- 2) Insert this SR The character set supported, and the maximum lengths of the <package identifier>, <class identifier>, <Java field name>, and <Java method name> are implementation-defined.
- 3) Insert after SR 18) Two <jar name>s are equivalent if and only if they have equivalent <jar id>s and equivalent implicit or explicit <schema name>s.

Access Rules

No additional Access Rules.

General Rules

- 1) Insert this GR A <jar name> identifies a JAR.
- 2) Insert this GR A <jar id> represents an unqualified JAR name.
- 3) Insert this GR A <Java class name> identifies a fully qualified Java class.
- 4) Insert this GR A <packages> identifies a fully qualified Java package.
- 5) Insert this GR A <package identifier> represents an unqualified Java package name.
- 6) Insert this GR A <class identifier> represents an unqualified Java class name.
- 7) Insert this GR A <Java field name> represents the name of a field within a Java class.
- 8) Insert this GR A <Java method name> represents the name of a method within a Java class.

Conformance Rules

No additional Conformance Rules.

6 Scalar expressions

This Clause modifies Clause 6, “Scalar expressions”, in ISO/IEC 9075-2.

6.1 <method invocation>

This Subclause modifies Subclause 6.16, “<method invocation>”, in ISO/IEC 9075-2.

Function

Reference an SQL-invoked method of a user-defined type value.

Format

No additional Format items.

Syntax Rules

- 1) Insert after SR 2) If UDT is an external Java data type, then <method invocation> shall immediately contain <direct invocation>.

Access Rules

No additional Access Rules.

General Rules

No additional General Rules.

Conformance Rules

No additional Conformance Rules.

6.2 <new specification>

This Subclause modifies Subclause 6.18, “<new specification>”, in ISO/IEC 9075-2.

Function

Invoke a method on a newly-constructed value of a structured type.

Format

No additional Format items.

Syntax Rules

No additional Syntax Rules.

Access Rules

No additional Access Rules.

General Rules

- 1) Insert this GR If Feature J571, “NEW operator” is not supported, then the mechanism used to invoke a constructor of an external Java data type is implementation-defined.

Conformance Rules

- 1) Insert this CR Without Feature J571, “NEW operator”, conforming SQL language shall not contain a <new specification> in which the schema identified by the implicit or explicit <schema name> of the <routine name> *RN* immediately contained in <routine invocation> immediately contained in the <new specification> contains a user-defined type whose user-defined type name is *RN* and that is an external Java data type.

7 Predicates

This Clause modifies [Clause 8, “Predicates”](#), in ISO/IEC 9075-2.

7.1 <comparison predicate>

This Subclause modifies [Subclause 8.2, “<comparison predicate>”](#), in ISO/IEC 9075-2.

Function

Specify a comparison of two row values.

Format

No additional Format items.

Syntax Rules

- 1) NOTE 7 — [Replace Note 179](#) The comparison form and comparison categories included in the user-defined type descriptors of both *UDT1* and *UDT2* are constrained to be the same and to be the same as those of all their supertypes. If the comparison category is COMPARABLE, then no comparison functions shall be specified for *T1* and *T2*. If the comparison category is either STATE or RELATIVE, then *UDT1* and *UDT2* are constrained to have the same comparison function; if the comparison category is MAP, they are not constrained to have the same comparison function.

Access Rules

No additional Access Rules.

General Rules

- 1) [Insert before GR 1\)b\)iv\)3\)](#) and its subrules If the comparison category of *UDT_x* is COMPARABLE, then:
 - a) The subject SQL data type shall be an external Java data type. Let *JC* be the subject Java class of that external Java data type.
 NOTE 8 — Syntax Rules in [Subclause 9.11, “<user-defined ordering definition>”](#), require that *JC* implement the Java interface `java.lang.Comparable`. The interface `java.lang.Comparable` requires an implementing Java class to have a method named `compareTo`, whose result data type is `JavaInt`.
 - b) Let *XJV* be the value of *X* in the associated JVM. Let *YJV* be the value of *Y* in that associated JVM.
 - c) $X = Y$

has the same result as if the JVM executed the Java boolean expression

`XJV.compareTo(YJV) == 0`

d) $X < Y$

has the same result as if the JVM executed the Java boolean expression

`XJV.compareTo(YJV) < 0`

e) $X <> Y$

has the same result as if the JVM executed the Java boolean expression

`XJV.compareTo(YJV) != 0`

f) $X > Y$

has the same result as if the JVM executed the Java boolean expression

`XJV.compareTo(YJV) > 0`

g) $X \leq Y$

has the same result as if the JVM executed the Java boolean expression

`XJV.compareTo(YJV) <= 0`

h) $X \geq Y$

has the same result as if the JVM executed the Java boolean expression

`XJV.compareTo(YJV) >= 0`

Conformance Rules

No additional Conformance Rules.

8 Additional common elements

This Clause modifies [Clause 10](#), “Additional common elements”, in ISO/IEC 9075-2.

8.1 <Java parameter declaration list>

Function

Specify the Java types of parameters for a Java method.

Format

<Java parameter declaration list> ::= <left paren> [<Java parameters>] <right paren>

<Java parameters> ::= <Java data type> [{ <comma> <Java data type> }...]

<Java data type> ::= !! See the Syntax Rules

Syntax Rules

- 1) A <Java data type> is a Java data type that is mappable or result set mappable, as specified in [Subclause 4.5](#), “Parameter mapping”. The <Java data type> names are case sensitive, and shall be fully qualified with their package names, if any.

Access Rules

None.

General Rules

None.

Conformance Rules

- 1) Without Feature J631, “Java signatures”, conforming SQL language shall not contain a <Java parameter declaration list> that is not equivalent to the default Java method signature as determined in [Subclause 8.6](#), “Java routine signature determination”.

8.2 <SQL Java path>

Function

Control the resolution of Java classes across installed JARs.

Format

<SQL Java path> ::= [<path element>...]

<path element> ::= <left paren> <referenced class> <comma> <resolution jar> <right paren>

<referenced class> ::=
 [<packages> <period>] <asterisk>
 | [<packages> <period>] <class identifier>

<resolution jar> ::= <jar name>

Syntax Rules

None.

Access Rules

None.

General Rules

- 1) When a Java class *CJ* in a JAR *J* is executed in an SQL-implementation, let *P* be the <SQL Java path> associated with *J* by an invocation of the SQLJ.ALTER_JAVA_PATH procedure.
- 2) Every static or dynamic reference in *CJ* to a class with the name *CN* that is not a system class and is not contained in *J* is resolved as follows.

For each <path element> *PE* (if any) in *P*, in the order in which they were specified:

- a) Let *RC* and *RJ* be the <referenced class> and <resolution jar>, respectively, contained in *PE*. Let *JR* be the JAR referenced by *RJ*.
- b) If *RJ* is not the name of an installed JAR, then an exception condition is raised: *Java execution — invalid JAR name in path*.

NOTE 9 — This exception can only occur if the implementation-defined action taken for an SQLJ.ALTER_JAVA_PATH call that raised an exception results in leaving invalid <jar name>s in the SQL-Java path.

- c) If *RC* is equivalent to *CN*, then:
 - i) If *CN* is the name of some class *C* in *JR*, then *CN* resolves to class *C*.

- ii) If *CN* is not the name of a class in *JR*, then an exception condition is raised: *Java execution — unresolved class name*.
- d) If *RC* simply contains <asterisk> and simply contains <packages>, then let *PKG* be the specified <packages> and let *CI* be the <class identifier> of *CN*. If the <Java class name> of *CN* is *PKG.CI*, then:
 - i) If *CN* is the name of a class *C* in *JR*, then *CN* resolves to class *C*.
 - ii) If *CN* is not the name of a class in *JR*, then an exception condition is raised: *Java execution — unresolved class name*.
- e) If *RC* simply contains <asterisk> and does not simply contain <packages>, then:
 - i) If *CN* is the name of a class *C* in *JR*, then *CN* resolves to class *C*.
 - ii) If *CN* is not the name of a class in *RJ*, then *CN* is not resolved by the <path element> being considered and the next <path element> in *P* is considered.
- 3) If *CN* is not resolved after all <path element>s in *P* have been considered, then an exception condition is raised: *Java execution — unresolved class name*.

Conformance Rules

- 1) Without Feature J601, “SQL-Java paths”, conforming SQL language shall not contain an <SQL Java path>.

8.3 <routine invocation>

This Subclause modifies *Subclause 10.4*, “<routine invocation>”, in ISO/IEC 9075-2.

Function

Invoke an SQL-invoked routine.

Format

No additional Format items.

Syntax Rules

- 1) Insert this SR If *SR* is an external Java routine, then:
 - a) No <SQL argument> immediately contained in <SQL argument list> shall immediately contain <generalized expression>.
 - b) If validation of the <Java parameter declaration list> has been implementation-defined to be performed by <routine invocation>, then the Syntax Rules of *Subclause 8.6*, “Java routine signature determination”, are applied with <routine invocation>, a method specification index of 0 (zero), and subject routine *SR*.

Access Rules

No additional Access Rules.

General Rules

- 1) Insert after GR 3)b)i) If *R* is an external Java routine, then let CPV_i be an implementation-defined non-null value of declared type T_i .
- 2) Insert before GR 4) If *R* is an external Java routine that is not a static field method, then let *P* be the *subject Java method* of *R*.
 NOTE 10 — The subject Java method of an external Java routine is defined in *Subclause 8.6*, “Java routine signature determination”.
- 3) Replace the first paragraph of GR 4) If *R* is an external routine that is not an external Java routine, then:
- 4) Replace the first paragraph of GR 5)d)ii) If *R* is not a static field method, then:
- 5) Insert before GR 8)d) If *R* specifies PARAMETER STYLE JAVA, then

Case:

- a) If *R* is an SQL-invoked function that is an array-returning external function or a multiset-returning external function, then the effective SQL parameter list *ESPL* of *R* is set as follows:

- i) If R 's returned array's element type or returned multiset's element type is a row type, then let FRN be the degree of the element type; otherwise, let FRN be 1 (one).
- ii) For i ranging from 1 (one) to PN , the i -th entry in $ESPL$ is set to CPV_i .
- iii) For i ranging from $PN+1$ to $PN+FRN$, the i -th entries in $ESPL$ are the *result data items*.
- iv) For i equal to $PN+FRN+1$, the i -th entry in $ESPL$ is the save area data item and for i equal to $PN+FRN+2$, the i -th entry in $ESPL$ is the *call type data item*.
- v) Set the value of the save area data item (that is, SQL argument value list entry $PN+FRN+1$) to null and set the value of the call type data item (that is, SQL argument value list entry $PN+FRN+2$) to -1 .

NOTE 11 — Initialization of the save area data item occurs in Subclause 8.5, “Execution of array-returning functions”; for now, it is set to null.

- b) Otherwise, for i ranging from 1 (one) to PN , let the effective SQL parameter list $ESPL$ of R be the list of values CPV_i .
- 6) Replace the first paragraph of GR 8)f)ii)1) If R is not an external Java routine and R is neither an array-returning external function nor a multiset-returning external function, then P is executed with a list of EN parameters PD_i whose parameter names are PN_i and whose values are set as follows:
- 7) Insert before GR 8)f)ii)2) If R is an external Java routine and R is not an array-returning external function or a multiset-returning external function, then P is executed in a manner determined as follows and with a list of parameters PD_i whose values are set as follows:
- a) Let SRD be routine descriptor of R .
 - b) If SRD indicates that R is an SQL-invoked method, then let $SRUDT$ be the user-defined type whose descriptor contains SR 's corresponding method specification descriptor MSD and let $JCLSN$ be the subject Java class of $SRUDT$.
 - c) Case:
 - i) If SRD indicates that R is an SQL-invoked method and MSD indicates that R is a static field method, then:
 - 1) Let JSF be the subject static field of R .

NOTE 12 — The “subject static field” of an SQL-invoked method is defined in Subclause 8.6, “Java routine signature determination”.
 - 2) Let ERT be the effective returns data type of R .

NOTE 13 — “effective returns data type” is defined in the Syntax Rules of Subclause 10.4, “<routine invocation>”, in ISO/IEC 9075-2.
 - 3) Case:
 - A) If ERT is a user-defined type, then
 - I) Let $SJCE$ be the most specific Java class of the value of JSF , and let STU be the user-defined type whose subject Java class is $SJCE$ and whose user-defined type is ERT or is a subclass of ERT .

- II) Let *UIS* be the <interface specification> specified by the user-defined type descriptor of *STU*.

Case:

- 1) If *UIS* is *SERIALIZABLE*, then:

- a) The subject Java class *SJCE*'s `writeObject()` method is executed to convert the Java value of *JSF* to the SQL value *SSFV* of user-defined type *STU*.
- b) The method of execution of the subject Java class's implementation of `writeObject()` is implementation-defined.

NOTE 14 — If *UIS* is *SERIALIZABLE*, then, as described in Subclause 9.4, “<user-defined type definition>”, the descriptor's subject Java class implements the Java interface `java.io.Serializable` and defines that interface's `writeObject()` method as described by the [J2SE].

- 2) If *UIS* is *SQLDATA*, then:

- a) The subject Java class *SJCE*'s method `writeSQL()` is executed to convert the Java value of *JSF* to the SQL value *SSFV* of user-defined type *STU*.
- b) The method of execution of the subject Java class's implementation of `writeSQL()` is implementation-defined.

NOTE 15 — If *UIS* is *SQLDATA*, then, as described in Subclause 9.4, “<user-defined type definition>”, the descriptor's subject Java class implements the Java interface `java.sql.SQLData` and defines that interface's `writeSQL()` method as described by [JDBC] and [J2SE].

B) Otherwise, the value of *SSFV* is set to the value of *JSF*.

- 4) Let the result of the <routine invocation> be the result of assigning *SSFV* to a target of declared type *ERT* according to the rules of Subclause 9.2, “Store assignment”, in ISO/IEC 9075-2. No further General Rules of this Subclause are applied.

ii) Otherwise:

- 1) Let *JPDL* be an ordered list of the data types of the Java parameters declared for *P* in the order they appear in *P*'s declaration.

NOTE 16 — If any Java parameter is declared to be of an array class, then *JPDL* reflects that information.

- 2) If *SRD* indicates that *R* is an SQL-invoked method and *MSD* indicates that *R* is an instance method or a constructor method, then prefix *JPDL* with the subject parameter as follows.

Case:

A) If *JPDL* contains one or more Java data types, then prefix *JPDL* with *JCLSN*.

B) Otherwise, replace *JPDL* with *JCLSN*.

- 3) Let *JP_i* be the *i*-th data type in *JPDL*.

- 4) For i ranging from 1 (one) to EN , if JP_i is of an array class, then let JP_i be the component type of JP_i .

NOTE 17 — The component type of a Java array is defined in [Java].

- 5) For i ranging from 1 (one) to EN , if ESP_i is the SQL null value and if JP_i is any of `boolean`, `byte`, `short`, `int`, `long`, `float`, or `double`, then an exception condition is raised: *external routine invocation exception — null value not allowed*.
- 6) For i ranging from 1 (one) to EN ,

Case:

- A) If the value of ESP_i is a user-defined type, then let the most specific type of ESP_i be U , let UIS be the <interface specification> specified by the user-defined type descriptor of U , and let $SJCU$ be the subject Java class of U .

Case:

- I) If UIS is `SERIALIZABLE`, then:

- 1) The subject Java class $SJCU$'s method `readObject()` is executed to convert the value of ESP_i to a Java object, the value of PD_i .
- 2) The method of execution of the subject Java class's implementation of `readObject()` is implementation-defined.

NOTE 18 — If UIS is `SERIALIZABLE`, then, as described in Subclause 9.4, “<user-defined type definition>”, the subject Java class of U implements the Java interface `java.io.Serializable` and defines that interface's `readObject()` method as described by [J2SE].

- II) If UIS is `SQLDATA`, then:

- 1) The subject Java class $SJCU$'s method `readSQL()` is executed to convert the value of ESP_i to a Java object, the value of PD_i .
- 2) The method of execution of the subject Java class's implementation of `readSQL()` is implementation-defined.

NOTE 19 — If UIS is `SQLDATA`, then, as described in Subclause 9.4, “<user-defined type definition>”, the subject Java class of U implements the Java interface `java.sql.SQLData` and defines that interface's `readSQL()` method as described by [JDBC] and [J2SE].

- B) Otherwise, the value of PD_i of the Java data type JP_i is set to the value of ESP_i .

- 7) For i ranging from 1 (one) to EN , if P_i is an output SQL parameter or both an input SQL parameter and an output SQL parameter, then:

- A) Let PAD_i be a Java array of length 1 (one) and data type JP_i initialized as specified in [Java].

NOTE 20 — PAD_i is a Java object effectively created by execution of the Java expression `new JP_i[1]`.

- B) If P_i is both an input SQL parameter and an output SQL parameter, then $PAD_i[0]$ is set to PD_i .

C) PD_i is replaced by PAD_i .

8) Let $JPEN$ be the number of Java data types in $JPDL$.

9) If $JPEN$ is greater than EN , then prepare the Java parameters for the DYNAMIC RESULT SET parameters as follows.

For i ranging from $EN+1$ to $JPEN$:

A) Let PAD_i be a Java array of length 1 (one) and data type JP_i initialized as specified in [Java].

NOTE 21 — PAD_i is a Java object effectively created by execution of the Java expression `new JP_i[1]`.

B) The value of PD_i is set to the value of PAD_i .

10) Let $JCLSN$, JMN , and ERT be respectively the subject Java class name, the subject Java method name, and the effective returns data type of R . The subject Java method of the subject Java class is invoked as follows.

Case:

A) If R is an SQL-invoked procedure, then:

I) If $JPEN$ is greater than 0 (zero), then the following Java statement is effectively executed:

$$JCLSN.JMN (PD_1, \dots, PD_{JPEN}) ;$$

II) If $JPEN$ equals 0 (zero), then the following Java statement is effectively executed:

$$JCLSN.JMN () ;$$

B) If R is an SQL-invoked method whose routine descriptor specifies STATIC or R is an SQL-invoked regular function, then:

I) If ERT is a user-defined type, then let $SJCE$ and $SJCEN$ be the subject Java class and the subject Java class name of ERT , respectively.

II) If ERT is not a user-defined type, then let $SJCEN$ be the Java returns data type of the subject Java method.

III) If $JPEN$ is greater than 0 (zero), then the following Java statement is effectively executed:

$$SJCEN \text{ tempU} = JCLSN.JMN (PD_1, \dots, PD_{JPEN}) ;$$

IV) If $JPEN$ equals 0 (zero), then the following Java statement is effectively executed:

$$SJCEN \text{ tempU} = JCLSN.JMN () ;$$

C) If R is an SQL-invoked constructor method, then: