

---

---

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

**Part 13:**

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

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

*Partie 13: Routines et types utilisant le langage de programmation Java™  
(SQL/JRT)*



**PDF disclaimer**

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

Adobe is a trademark of Adobe Systems Incorporated.

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

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

© ISO/IEC 2002

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

ISO copyright office  
Case postale 56 • CH-1211 Geneva 20  
Tel. + 41 22 749 01 11  
Fax + 41 22 749 09 47  
E-mail [copyright@iso.ch](mailto:copyright@iso.ch)  
Web [www.iso.ch](http://www.iso.ch)

Printed in Switzerland

## Contents

## Page

Foreword. . . . .	viii
Introduction. . . . .	ix
<b>1 Scope. . . . .</b>	<b>1</b>
<b>2 Normative references . . . . .</b>	<b>3</b>
2.1 JTC1 standards . . . . .	3
2.2 Publicly-available specifications . . . . .	3
<b>3 Definitions, notations, and conventions . . . . .</b>	<b>5</b>
3.1 Definitions . . . . .	5
3.1.1 Definitions provided in Part 13 . . . . .	5
3.1.2 Definitions taken from Java . . . . .	6
3.2 Notations . . . . .	6
3.3 Conventions . . . . .	7
3.3.1 Specification of built-in procedures . . . . .	7
3.3.2 Specification of deployment descriptor files . . . . .	7
3.3.3 Relationships to other parts of ISO/IEC 9075 . . . . .	8
3.3.3.1 Clause, Subclause, and Table relationships . . . . .	8
3.4 Object identifier for Database Language SQL . . . . .	12
<b>4 Concepts. . . . .</b>	<b>17</b>
4.1 The Java programming language . . . . .	17
4.2 SQL-invoked routines . . . . .	18
4.3 Java class name resolution . . . . .	20
4.4 SQL result sets . . . . .	20
4.5 Parameter mapping . . . . .	21
4.6 Unhandled Java exceptions . . . . .	22
4.7 Data types . . . . .	23
4.8 User-defined types . . . . .	23
4.8.1 User-defined type comparison and assignment . . . . .	26
4.8.2 Accessing static fields . . . . .	26
4.8.3 Converting objects between SQL and Java . . . . .	27
4.8.3.1 SERIALIZABLE . . . . .	27
4.8.3.2 SQLDATA . . . . .	27
4.8.3.3 Developing for portability . . . . .	28

4.9	Built-in procedures	28
4.10	Privileges	29
4.11	JARs	29
4.11.1	Deployment descriptor files	30
<b>5</b>	<b>Lexical elements</b>	<b>31</b>
5.1	<token> and <separator>	31
5.2	Names and identifiers	32
<b>6</b>	<b>Scalar expressions</b>	<b>35</b>
6.1	<method invocation>	35
6.2	<new specification>	36
<b>7</b>	<b>Predicates</b>	<b>37</b>
7.1	<comparison predicate>	37
<b>8</b>	<b>Additional common elements</b>	<b>39</b>
8.1	<Java parameter declaration list>	39
8.2	<SQL Java path>	40
8.3	<routine invocation>	42
8.4	<language clause>	51
8.5	Java routine signature determination	52
<b>9</b>	<b>Schema definition and manipulation</b>	<b>59</b>
9.1	<table definition>	59
9.2	<view definition>	60
9.3	<user-defined type definition>	61
9.4	<attribute definition>	65
9.5	<alter type statement>	69
9.6	<drop data type statement>	70
9.7	<SQL-invoked routine>	71
9.8	<alter routine statement>	74
9.9	<drop routine statement>	75
9.10	<user-defined ordering definition>	76
9.11	<drop user-defined ordering statement>	78
<b>10</b>	<b>Access control</b>	<b>79</b>
10.1	<grant privilege statement>	79
10.2	<privileges>	80
10.3	<revoke statement>	81
<b>11</b>	<b>Built-in procedures</b>	<b>83</b>
11.1	SQLJ.INSTALL_JAR procedure	83
11.2	SQLJ.REPLACE_JAR procedure	85
11.3	SQLJ.REMOVE_JAR procedure	87
11.4	SQLJ.ALTER_JAVA_PATH procedure	89

<b>12</b>	<b>Java topics</b>	91
12.1	Java facilities supported by this part of ISO/IEC 9075	91
12.1.1	Package java.sql	91
12.1.2	System properties	91
12.2	Deployment descriptor files	92
<b>13</b>	<b>Information Schema</b>	95
13.1	JAR_JAR_USAGE view	95
13.2	JARS view	96
13.3	METHOD_SPECIFICATIONS view	97
13.4	ROUTINE_JAR_USAGE view	98
13.5	TYPE_JAR_USAGE view	99
13.6	USER_DEFINED_TYPES view	100
13.7	Short name views	101
<b>14</b>	<b>Definition Schema</b>	103
14.1	JAR_JAR_USAGE base table	103
14.2	JARS base table	104
14.3	METHOD_SPECIFICATIONS base table	105
14.4	ROUTINE_JAR_USAGE base table	106
14.5	ROUTINES base table	107
14.6	TYPE_JAR_USAGE base table	108
14.7	USAGE_PRIVILEGES base table	109
14.8	USER_DEFINED_TYPES base table	110
<b>15</b>	<b>Status codes</b>	113
15.1	Class and subclass values for uncaught Java exceptions	113
15.2	SQLSTATE	114
<b>16</b>	<b>Conformance</b>	115
16.1	Claims of conformance	115
<b>Annex A</b>	<b>SQL Conformance Summary</b>	117
<b>Annex B</b>	<b>Implementation-defined elements</b>	123
<b>Annex C</b>	<b>Implementation-dependent elements</b>	127
<b>Annex D</b>	<b>SQL Feature Taxonomy</b>	129
<b>Annex E</b>	<b>Routines tutorial</b>	131
E.1	Technical components	131
E.2	Overview	132
E.3	Example Java methods: region and correctStates	133
E.4	Installing region and correctStates in SQL	133
E.5	Defining SQL names for region and correctStates	135
E.6	A Java method with output parameters: bestTwoEmps	136
E.7	A CREATE PROCEDURE best2 for bestTwoEmps	137

E.8	Calling the best2 procedure .....	138
E.9	A Java method returning a result set: orderedEmps .....	138
E.10	A CREATE PROCEDURE rankedEmps for orderedEmps .....	140
E.11	Calling the rankedEmps procedure .....	141
E.12	Overloading Java method names and SQL names .....	141
E.13	Java main methods .....	143
E.14	Java method signatures in the CREATE statements .....	144
E.15	Null argument values and the RETURNS NULL clause .....	145
E.16	Static variables .....	147
E.17	Dropping SQL names of Java methods .....	148
E.18	Removing Java classes from SQL .....	148
E.19	Replacing Java classes in SQL .....	149
E.20	Visibility .....	150
E.21	Exceptions .....	150
E.22	Deployment descriptors .....	151
E.23	Paths .....	154
E.24	Privileges .....	156
E.25	Information Schema .....	156
<b>Annex F</b>	<b>Types tutorial .....</b>	<b>157</b>
F.1	Overview .....	157
F.2	Example Java classes .....	157
F.3	Installing Address and Address2Line in an SQL system .....	159
F.4	CREATE TYPE for Address and Address2Line .....	160
F.5	Multiple SQL types for a single Java class .....	162
F.6	Collapsing subclasses .....	162
F.7	GRANT and REVOKE statements for data types .....	164
F.8	Deployment descriptors for classes .....	164
F.9	Using Java classes as data types .....	166
F.10	SELECT, INSERT, and UPDATE .....	166
F.11	Referencing Java fields and methods in SQL .....	167
F.12	Extended visibility rules .....	168
F.13	Logical representation of Java instances in SQL .....	168
F.14	Static methods .....	170
F.15	Static fields .....	170
F.16	Instance-update methods .....	171
F.17	Subtypes in SQL/JRT data .....	173
F.18	References to fields and methods of null instances .....	174
F.19	Ordering of SQL/JRT data .....	175
<b>Annex G</b>	<b>Incompatibilities with ANSI NCITS 331 .....</b>	<b>177</b>
G.1	References .....	177
G.2	Incompatibilities .....	177
<b>Index</b>	<b>.....</b>	<b>181</b>

TABLES

Tables		Page
1	Clause, Subclause, and Table relationships .....	8
2	System properties .....	91
3	SQLSTATE class and subclass values .....	114
4	Feature taxonomy for features outside Core SQL .....	129

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 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. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 3.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this part of ISO/IEC 9075 may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 9075-13 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 32, *Data management and interchange*.

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

- *Part 1: Framework (SQL/Framework)*
- *Part 2: Foundation (SQL/Foundation)*
- *Part 3: Call-Level Interface (SQL/CLI)*
- *Part 4: Persistent Stored Modules (SQL/PSM)*
- *Part 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<sup>TM</sup> Programming Language (SQL/JRT)*

Annexes A, B, C, D, E, F and G of this part of ISO/IEC 9075 are for information only.



## 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 to Core SQL 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.
- 23) Annex G, "Incompatibilities with ANSI NCITS 331", is an informative Annex. It lists the incompatibilities between this edition of this part of ISO/IEC 9075 and NCITS 331.1 and NCITS 331.2.

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.

## 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 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.)

(Blank page)

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 2 Normative references

The following normative documents contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 9075. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. However, parties to agreements based on this part of ISO/IEC 9075 are encouraged to investigate the possibility of applying the most recent editions of the normative documents indicated below. For undated references, the latest edition of the normative document referred to applies. Members of ISO and IEC maintain registers of currently valid International Standards.

### 2.1 JTC 1 standards

ISO 8824-1:1998, *Information technology — Abstract Syntax Notation One (ASN.1): Specification of basic notation*

ISO/IEC 9075-1:1999, *Information technology — Database languages — SQL — Part 1: Framework (SQL/Framework)*

ISO/IEC 9075-2:1999, *Information technology — Database languages — SQL — Part 2: Foundation (SQL/Foundation)*

ISO/IEC 9075-5:1999, *Information technology — Database languages — SQL — Part 5: Host Language Bindings (SQL/Bindings)*

ISO/IEC 9075-10:2000, *Information technology — Database languages — SQL — Part 10: Object Language Bindings (SQL/OLB)*

### 2.2 Publicly-available specifications

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

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

*Java 2 Platform, Standard Edition, v1.2.2, API Specification*,  
<http://web2.java.sun.com/products/jdk/1.2/docs/api/>.

*Java Object Serialization Specification*,  
<http://web2.java.sun.com/products/jdk/1.2/docs/guide/serialization/spec/serialTOC.doc.html>.

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

**2.2 Publicly-available specifications**

*JDBC™ 2.0 API, Version 1.0*, Seth White & Mark Hapner, Sun Microsystems, Inc., 30 May, 1999.

*JDBC 2.0 Standard Extension API, Version 1.0*, Seth White & Mark Hapner, Sun Microsystems, Inc., 7 December, 1998.

*JDBC API Tutorial and Reference, Second Edition: Universal Data Access for the Java 2 Platform*, Seth White, Maydene Fisher, Rick Cattell, Graham Hamilton, and Mark Hapner, Addison Wesley, Reading MA, 1999, ISBN 0-201-43328-1.

*W3C Architecture domain: Naming and Addressing (URLs)*,  
<http://www.w3.org/Addressing/Activity.html>.

RFC 1738, *Uniform Resource Locators (URL)*, T. Berners-Lee, L. Masinter, M. McCahill, December, 1994.

RFC 1808, *Relative Uniform Resource Locators*, R. Fielding, June, 1995.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

### 3 Definitions, notations, and conventions

#### 3.1 Definitions

##### 3.1.1 Definitions provided in Part 13

Insert this paragraph For the purposes of this part of ISO/IEC 9075, the definitions given in ISO/IEC 9075-1, ISO/IEC 9075-2, ISO/IEC 9075-5, and ISO/IEC 9075-10, and the following definitions, apply.

- a) **class file:** A file containing the compiled byte code for a Java class.
- b) **default connection:** a JDBC connection to the current SQL-implementation, SQL-session, and SQL-transaction established with the data source URL 'jdbc:default:connection'.
- c) **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*. For example, when a JAR is installed, one or more <SQL-invoked routine>s that specify LANGUAGE JAVA and either PROCEDURE or FUNCTION and the associated <grant privilege statement>s can be specified in the deployment descriptor and executed as part of the install process.
- d) **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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*, specifies SQLJDeploymentDescriptor: TRUE.
- e) **external Java data type:** an SQL user-defined type defined with a <user-defined type definition> that specifies an <external Java type clause>.
- f) **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>.
- g) **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.
- h) **Java Archive (JAR):** a zip formatted file, as described by the java.util.zip section of *Java 2 Platform, Standard Edition, v1.2.2, API Specification*, 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.
- i) **JVM:** A Java Virtual Machine, as defined by *The Java Virtual Machine Specification, Second Edition*.
- j) **ser file:** A file containing representations of Java objects in the form defined in *Java Object Serialization Specification*.

### 3.1 Definitions

- k) **subject Java class**: the Java class uniquely identified by the combination of the class's *subject Java class name* and its containing JAR.
- l) **subject Java class name**: the fully-qualified package and class name of a Java class.
- m) **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.

#### 3.1.2 Definitions taken from Java

This part of ISO/IEC 9075 makes use of the following terms defined in *The Java Language Specification, Second Edition*:

- a) **block**
- b) **class declaration**
- c) **class instance**
- d) **class variable**
- e) **field**
- f) **instance initializer**
- g) **instance variable**
- h) **interface**
- i) **local variable**
- j) **nested class**
- k) **package**
- l) **static initializer**
- m) **subpackage**

This part of ISO/IEC 9075 makes use of the following terms defined in *The Java Virtual Machine Specification, Second Edition*:

- a) **class file**
- b) **Java Virtual Machine**

### 3.2 Notations

Insert this paragraph The syntax notation used in this part of ISO/IEC 9075 is an extended version of BNF ("Backus Normal Form" or "Backus Naur Form"). This version of BNF is fully described in Subclause 6.1, "Notation", of ISO/IEC 9075-1.



## 3.3 Conventions

Insert this paragraph Except as otherwise specified in this part of ISO/IEC 9075, the conventions used in this part of ISO/IEC 9075 are identical to those described in ISO/IEC 9075-1 and ISO/IEC 9075-2.

### 3.3.1 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 must 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 must 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.

### 3.3.2 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.

## 3.3 Conventions

## 3.3.3 Relationships to other parts of ISO/IEC 9075

## 3.3.3.1 Clause, Subclause, and Table relationships

Table 1—Clause, Subclause, and Table relationships

Clause, Subclause, or Table in this part of ISO/IEC 9075	Corresponding Clause, Subclause, or Table from another part	Part containing correspondence
Clause 3, "Definitions, notations, and conventions"	Clause 3, "Definitions, notations, and conventions"	ISO/IEC 9075-2
Subclause 3.1, "Definitions"	Subclause 3.1, "Definitions"	ISO/IEC 9075-2
Subclause 3.1.1, "Definitions provided in Part 13"	(none)	(none)
Subclause 3.1.2, "Definitions taken from Java"	(none)	(none)
Subclause 3.2, "Notations"	Subclause 3.2, "Notation"	ISO/IEC 9075-2
Subclause 3.3, "Conventions"	Subclause 3.3, "Conventions"	ISO/IEC 9075-2
Subclause 3.3.1, "Specification of built-in procedures"	(none)	(none)
Subclause 3.3.2, "Specification of deployment descriptor files"	(none)	(none)
Subclause 3.3.3, "Relationships to other parts of ISO/IEC 9075"	(none)	(none)
Subclause 3.3.3.1, "Clause, Subclause, and Table relationships"	(none)	(none)
Subclause 3.4, "Object identifier for Database Language SQL"	Subclause 6.3, "Object identifier for Database Language SQL"	ISO/IEC 9075-1
Clause 4, "Concepts"	Clause 4, "Concepts"	ISO/IEC 9075-2
Subclause 4.1, "The Java programming language"	(None)	(None)
Subclause 4.2, "SQL-invoked routines"	Subclause 4.23, "SQL-invoked routines"	ISO/IEC 9075-2
Subclause 4.3, "Java class name resolution"	(None)	(None)
Subclause 4.4, "SQL result sets"	(None)	(None)
Subclause 4.5, "Parameter mapping"	(None)	(None)
Subclause 4.6, "Unhandled Java exceptions"	(None)	(None)
Subclause 4.7, "Data types"	Subclause 4.1, "Data types"	ISO/IEC 9075-2
Subclause 4.8, "User-defined types"	Subclause 4.8, "User-defined types"	ISO/IEC 9075-2

Table 1—Clause, Subclause, and Table relationships (Cont.)

Clause, Subclause, or Table in this part of ISO/IEC 9075	Corresponding Clause, Subclause, or Table from another part	Part containing correspondence
Subclause 4.8.1, “User-defined type comparison and assignment”	Subclause 4.8.4, "User-defined type comparison and assignment"	ISO/IEC 9075-2
Subclause 4.8.2, “Accessing static fields”	(None)	(None)
Subclause 4.8.3, “Converting objects between SQL and Java”	(None)	(None)
Subclause 4.8.3.1, “SERIALIZABLE”	(None)	(None)
Subclause 4.8.3.2, “SQLDATA”	(None)	(None)
Subclause 4.8.3.3, “Developing for portability”	(None)	(None)
Subclause 4.9, “Built-in procedures”	(None)	(None)
Subclause 4.10, “Privileges”	Subclause 4.31.2, "Privileges"	ISO/IEC 9075-2
Subclause 4.11, “JARs”	(None)	(None)
Subclause 4.11.1, “Deployment descriptor files”	(None)	(None)
Clause 5, “Lexical elements”	Clause 5, "Lexical elements"	ISO/IEC 9075-2
Subclause 5.1, “<token> and <separator>”	Subclause 5.1, "<token> and <separator>"	ISO/IEC 9075-2
Subclause 5.2, “Names and identifiers”	Subclause 5.1, "<token> and <separator>"	ISO/IEC 9075-2
Clause 6, “Scalar expressions”	Clause 6, "Scalar Expressions"	ISO/IEC 9075-2
Subclause 6.1, “<method invocation>”	Subclause 6.11, "<method invocation>"	ISO/IEC 9075-2
Subclause 6.2, “<new specification>”	Subclause 6.24, "<new specification>"	ISO/IEC 9075-2
Clause 7, “Predicates”	Clause 8, "Predicates"	ISO/IEC 9075-2
Subclause 7.1, “<comparison predicate>”	Subclause 8.2, "<comparison predicate>"	ISO/IEC 9075-2
Clause 8, “Additional common elements”	Clause 10, "Additional common elements"	ISO/IEC 9075-2
Subclause 8.1, “<Java parameter declaration list>”	(None)	(None)
Subclause 8.2, “<SQL Java path>”	(None)	(None)
Subclause 8.3, “<routine invocation>”	Subclause 10.4, "<routine invocation>"	ISO/IEC 9075-2
Subclause 8.4, “<language clause>”	Subclause 10.2, "<language clause>"	ISO/IEC 9075-2

**Table 1—Clause, Subclause, and Table relationships (Cont.)**

<b>Clause, Subclause, or Table in this part of ISO/IEC 9075</b>	<b>Corresponding Clause, Subclause, or Table from another part</b>	<b>Part containing correspondence</b>
Subclause 8.5, “Java routine signature determination”	(None)	(None)
Clause 9, “Schema definition and manipulation”	Clause 11, "Schema definition and manipulation"	ISO/IEC 9075-2
Subclause 9.1, “<table definition>”	Subclause 11.3, "<table definition>"	ISO/IEC 9075-2
Subclause 9.2, “<view definition>”	Subclause 11.21, "<view definition>"	ISO/IEC 9075-2
Subclause 9.3, “<user-defined type definition>”	Subclause 11.40, "<user-defined type definition>"	ISO/IEC 9075-2
Subclause 9.4, “<attribute definition>”	Subclause 11.41, "<attribute definition>"	ISO/IEC 9075-2
Subclause 9.5, “<alter type statement>”	Subclause 11.42, "<alter type statement>"	ISO/IEC 9075-2
Subclause 9.6, “<drop data type statement>”	Subclause 11.48, "<drop data type statement>"	ISO/IEC 9075-2
Subclause 9.7, “<SQL-invoked routine>”	Subclause 11.49, "<SQL-invoked routine>"	ISO/IEC 9075-2
Subclause 9.8, “<alter routine statement>”	Subclause 11.50, "<alter routine statement>"	ISO/IEC 9075-2
Subclause 9.9, “<drop routine statement>”	Subclause 11.51, "<drop routine statement>"	ISO/IEC 9075-2
Subclause 9.10, “<user-defined ordering definition>”	Subclause 11.54, "<user-defined ordering definition>"	ISO/IEC 9075-2
Subclause 9.11, “<drop user-defined ordering statement>”	Subclause 11.55, "<drop user-defined ordering statement>"	ISO/IEC 9075-2
Clause 10, “Access control”	Clause 12, "Access control"	ISO/IEC 9075-2
Subclause 10.1, “<grant privilege statement>”	Subclause 12.2, "<grant privilege statement>"	ISO/IEC 9075-2
Subclause 10.2, “<privileges>”	Subclause 10.5, "<privileges>"	ISO/IEC 9075-2
Subclause 10.3, “<revoke statement>”	Subclause 12.6, "<revoke statement>"	ISO/IEC 9075-2
Clause 11, “Built-in procedures”	(None)	(None)
Subclause 11.1, “SQLJ.INSTALL_JAR procedure”	(None)	(None)
Subclause 11.2, “SQLJ.REPLACE_JAR procedure”	(None)	(None)
Subclause 11.3, “SQLJ.REMOVE_JAR procedure”	(None)	(None)

Table 1—Clause, Subclause, and Table relationships (Cont.)

Clause, Subclause, or Table in this part of ISO/IEC 9075	Corresponding Clause, Subclause, or Table from another part	Part containing correspondence
Subclause 11.4, “SQLJ.ALTER_JAVA_PATH procedure”	(None)	(None)
Clause 12, “Java topics”	(None)	(None)
Subclause 12.1, “Java facilities supported by this part of ISO/IEC 9075”	(None)	(None)
Subclause 12.1.1, “Package java.sql”	(None)	(None)
Subclause 12.1.2, “System properties”	(None)	(None)
Subclause 12.2, “Deployment descriptor files”	(None)	(None)
Clause 13, “Information Schema”	Clause 20, "Information Schema"	ISO/IEC 9075-2
Subclause 13.1, “JAR_JAR_USAGE view”	(None)	(None)
Subclause 13.2, “JARS view”	(None)	(None)
Subclause 13.3, “METHOD_SPECIFICATIONS view”	Subclause 20.32, "METHOD_SPECIFICATIONS view"	ISO/IEC 9075-2
Subclause 13.4, “ROUTINE_JAR_USAGE view”	(None)	(None)
Subclause 13.5, “TYPE_JAR_USAGE view”	(None)	(None)
Subclause 13.6, “USER_DEFINED_TYPES view”	Subclause 20.65, "USER_DEFINED_TYPES view"	ISO/IEC 9075-2
Subclause 13.7, “Short name views”	Subclause 20.69, "Short name views"	ISO/IEC 9075-2
Clause 14, “Definition Schema”	Clause 21, "Definition Schema"	ISO/IEC 9075-2
Subclause 14.1, “JAR_JAR_USAGE base table”	(None)	(None)
Subclause 14.2, “JARS base table”	(None)	(None)
Subclause 14.3, “METHOD_SPECIFICATIONS base table”	Subclause 21.24, "METHOD_SPECIFICATIONS base table"	ISO/IEC 9075-2
Subclause 14.4, “ROUTINE_JAR_USAGE base table”	(None)	(None)
Subclause 14.5, “ROUTINES base table”	Subclause 21.33, "ROUTINES base table"	ISO/IEC 9075-2
Subclause 14.6, “TYPE_JAR_USAGE base table”	(None)	(None)

**Table 1—Clause, Subclause, and Table relationships (Cont.)**

<b>Clause, Subclause, or Table in this part of ISO/IEC 9075</b>	<b>Corresponding Clause, Subclause, or Table from another part</b>	<b>Part containing correspondence</b>
Subclause 14.7, "USAGE_PRIVILEGES base table"	Subclause 21.50, "USAGE_PRIVILEGES base table"	ISO/IEC 9075-2
Subclause 14.8, "USER_DEFINED_TYPES base table"	Subclause 21.52, "USER_DEFINED_TYPES base table"	ISO/IEC 9075-2
Clause 15, "Status codes"	Clause 22, "Status codes"	ISO/IEC 9075-2
Subclause 15.1, "Class and subclass values for uncaught Java exceptions"	(None)	(None)
Subclause 15.2, "SQLSTATE"	Subclause 22.1, "SQLSTATE"	ISO/IEC 9075-2
Clause 16, "Conformance"	Clause 8, "Conformance"	ISO/IEC 9075-1
Subclause 16.1, "Claims of conformance"	Subclause 8.2.3, "Claims of conformance"	ISO/IEC 9075-1
Annex A, "SQL Conformance Summary"	Annex A, "SQL Conformance Summary"	ISO/IEC 9075-2
Annex B, "Implementation-defined elements"	Annex B, "Implementation-defined elements"	ISO/IEC 9075-2
Annex C, "Implementation-dependent elements"	Annex C, "Implementation-dependent elements"	ISO/IEC 9075-2
Annex D, "SQL Feature Taxonomy"	Annex F, "SQL feature and package taxonomy"	ISO/IEC 9075-2
Annex E, "Routines tutorial"	(None)	(None)
Annex F, "Types tutorial"	(None)	(None)
Annex G, "Incompatibilities with ANSI NCITS 331"	(None)	(None)
Table 1, "Clause, Subclause, and Table relationships"	(none)	(none)
Table 2, "System properties"	(None)	(None)
Table 3, "SQLSTATE class and subclass values"	Table 27, "SQLSTATE class and subclass values"	ISO/IEC 9075-2
Table 4, "Feature taxonomy for features outside Core SQL"	Table 32, "SQL/Foundation feature taxonomy for features outside Core SQL"	ISO/IEC 9075-2

## 3.4 Object identifier for Database Language SQL

The object identifier for Database Language SQL is defined in Subclause 6.3, "Object identifier for Database Language SQL", of ISO/IEC 9075-1 with the following additions:

## Format

```
<invoked routine languages> ::=
    !! All alternatives from ISO/IEC 9075-1
    | <invoked Java>

<invoked Java> ::=
    8 | invokedJava <left paren> 8 <right paren>

<Part 13 yes> ::=
    <Part 13 conformance>
    <Part 13 routines>
    <Part 13 types>
    <Part 13 JAR privileges>
    <Part 13 interfaces>

<Part 13 conformance> ::=
    13
    | sqljrt2002 <left paren> 13 <right paren>

<Part 13 routines> ::=
    <Part 13 routines no>
    | <Part 13 routines yes> <routines support>

<Part 13 routines no> ::=
    0 | routinesno <left paren> 0 <right paren>

<Part 13 routines yes> ::=
    1 | routinesyes <left paren> 1 <right paren>

<Part 13 types> ::=
    <Part 13 types no>
    | <Part 13 types yes> <types support>

<Part 13 types no> ::=
    0 | typesno <left paren> 0 <right paren>

<Part 13 types yes> ::=
    1 | typesyes <left paren> 1 <right paren>

<Part 13 JAR privileges> ::=
    <Part 13 JAR privileges no>
    | <Part 13 JAR privileges yes> <JAR privileges support>

<Part 13 JAR privileges no> ::=
    0 | JARprivilegesno <left paren> 0 <right paren>

<Part 13 JAR privileges yes> ::=
    1 | JARprivilegesyes <left paren> 1 <right paren>

<Part 13 interfaces> ::=
    <Part 13 serializable> <Part 13 sqldata>

<Part 13 serializable> ::=
    <Part 13 serializable no>
    | <Part 13 serializable yes>

<Part 13 serializable no> ::=
    0 | serializableno <left paren> 0 <right paren>
```

### 3.4 Object identifier for Database Language SQL

```

<Part 13 serializable yes> ::=
    1 | serializableyes <left paren> 1 <right paren>

<Part 13 sqldata> ::=
    <Part 13 sqldata no>
    | <Part 13 sqldata yes>

<Part 13 sqldata no> ::=
    0 | sqldatano <left paren> 0 <right paren>

<Part 13 sqldata yes> ::=
    1 | sqldatayes <left paren> 1 <right paren>

<routines support> ::=
    <routines commands> <routines deployment>

<routines commands> ::=
    <routines commands no>
    | <routines commands yes>

<routines commands no> ::=
    0 | routinescommandsno <left paren> 0 <right paren>

<routines commands yes> ::=
    1 | routinescommandsyas <left paren> 1 <right paren>

<routines deployment> ::=
    <routines deployment no>
    | <routines deployment yes>

<routines deployment no> ::=
    0 | routinesdeploymentno <left paren> 0 <right paren>

<routines deployment yes> ::=
    1 | routinesdeploymentyes <left paren> 1 <right paren>

<types support> ::=
    <types commands> <types deployment>

<types commands> ::=
    <types commands no>
    | <types commands yes>

<types commands no> ::=
    0 | typescommandsnos <left paren> 0 <right paren>

<types commands yes> ::=
    1 | typescommandsyas <left paren> 1 <right paren>

<types deployment> ::=
    <types deployment no>
    | <types deployment yes>

<types deployment no> ::=
    0 | typesdeploymentno <left paren> 0 <right paren>

<types deployment yes> ::=
    1 | typesdeploymentyes <left paren> 1 <right paren>

<JAR privileges support> ::=

```



```

<JAR privileges commands> <JAR privileges deployment>

<JAR privileges commands> ::=
    <JAR privileges commands no>
    | <JAR privileges commands yes>

<JAR privileges commands no> ::=
    0 | JARprivilegescommand sno <left paren> 0 <right paren>

<JAR privileges commands yes> ::=
    1 | JARprivilegescommand syes <left paren> 1 <right paren>

<JAR privileges deployment> ::=
    <JAR privileges deployment no>
    | <JAR privileges deployment yes>

<JAR privileges deployment no> ::=
    0 | JARprivilegesdeploymentno <left paren> 0 <right paren>

<JAR privileges deployment yes> ::=
    1 | JARprivilegesdeploymenttyes <left paren> 1 <right paren>

```

## Syntax Rules

- 1) Insert this SR Specification of <Part 13 yes> implies that conformance to ISO/IEC 9075-13 is claimed.
- 2) Insert this SR If <Part 13 yes> is not specified then <invoked Java> shall not be specified.
- 3) Insert this SR If <Part 13 conformance> specifies `sqljrt2002`, then <SQL edition> shall specify <1999>.
- 4) Insert this SR Specification of <Part 13 routines no> implies that conformance to Feature J621, “external Java routines”, is not claimed.
- 5) Insert this SR Specification of <Part 13 routines yes> implies that conformance to Feature J621, “external Java routines”, is claimed.
- 6) Insert this SR Specification of <Part 13 types no> implies that conformance to Feature J622, “external Java types”, is not claimed.
- 7) Insert this SR Specification of <Part 13 types yes> implies that conformance to Feature J622, “external Java types”, is claimed.
- 8) Insert this SR Specification of <Part 13 JAR privileges no> implies that conformance to Feature J561, “JAR privileges”, is not claimed.
- 9) Insert this SR Specification of <Part 13 JAR privileges yes> implies that conformance to Feature J561, “JAR privileges”, is claimed.
- 10) Insert this SR Specification of <Part 13 serializable no> implies that conformance to Feature J541, “Serializable”, is not claimed.
- 11) Insert this SR Specification of <Part 13 serializable yes> implies that conformance to Feature J541, “Serializable”, is claimed.

### 3.4 Object identifier for Database Language SQL

- 12) Insert this SR Specification of <routines commands no> implies that conformance to Feature J511, “Commands”, is not claimed for Feature J621, “external Java routines”.
- 13) Insert this SR Specification of <routines commands yes> implies that conformance to Feature J511, “Commands”, is claimed for Feature J621, “external Java routines”.
- 14) Insert this SR Specification of <types commands no> implies that conformance to Feature J511, “Commands”, is not claimed for Feature J622, “external Java types”.
- 15) Insert this SR Specification of <types commands yes> implies that conformance to Feature J511, “Commands”, is claimed for Feature J622, “external Java types”.
- 16) Insert this SR Specification of <JAR privileges commands no> implies that conformance to Feature J511, “Commands”, is not claimed for Feature J561, “JAR privileges”.
- 17) Insert this SR Specification of <JAR privileges commands yes> implies that conformance to Feature J511, “Commands”, is claimed for Feature J561, “JAR privileges”.
- 18) Insert this SR Specification of <routines deployment no> implies that conformance to Feature J531, “Deployment”, is not claimed for Feature J621, “external Java routines”.
- 19) Insert this SR Specification of <routines deployment yes> implies that conformance to Feature J531, “Deployment”, is claimed for Feature J621, “external Java routines”.
- 20) Insert this SR Specification of <types deployment no> implies that conformance to Feature J531, “Deployment”, is not claimed for Feature J622, “external Java types”.
- 21) Insert this SR Specification of <types deployment yes> implies that conformance to Feature J531, “Deployment”, is claimed for Feature J622, “external Java types”.
- 22) Insert this SR Specification of <JAR privileges deployment no> implies that conformance to Feature J531, “Deployment”, is not claimed for Feature J561, “JAR privileges”.
- 23) Insert this SR Specification of <JAR privileges deployment yes> implies that conformance to Feature J531, “Deployment”, is claimed for Feature J561, “JAR privileges”.

## 4 Concepts

### 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.

## 4.1 The Java programming language

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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*, 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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*, 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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*, 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

Insert after 1st paragraph An SQL-invoked routine can be an SQL routine or an external routine. An SQL routine is an SQL-invoked routine whose <language clause> specifies SQL. The <routine body> of an SQL routine is an <SQL procedure statement>; the <SQL procedure statement> forming the <routine body> can be any SQL-statement, including an <SQL control statement>, but excluding an <SQL schema statement>, <SQL connection statement>, or <SQL transaction statement>.

Insert after 1st paragraph An external routine is one whose <language clause> does not specify SQL. The <routine body> of an external routine is an <external body reference> whose <external routine name> identifies a program written in some standard programming language other than SQL.

Insert after 1st 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 lead-in paragraph of the 10th paragraph A static SQL-invoked method, whether or not it is an external Java routine, satisfies the following conditions:

Delete the 13th and 14th paragraphs

Replace the 16th 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 18th 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 associate 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 must 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 must 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 by JDBC in *JDBC API Tutorial and Reference, Second Edition: Universal Data Access for the Java 2 Platform*, 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 must specify WITH RETURN, and the cursor must 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 must 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 must be of a class that implements the interface `java.sql.ResultSet` or the interface `sqlj.runtime.ResultSetIterator`, the JDBC ResultSet or SQLJ Iterator must have been explicitly assigned to the first element of that array, and that JDBC ResultSet or SQLJ Iterator must 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 *JDBC data type mapping table*, Table 47.1, entitled "JDBC Types mapped to Java Types", in *JDBC API Tutorial and Reference, Second Edition: Universal Data Access for the Java 2 Platform*. 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 *JDBC object type mapping table*, Table 47.3, entitled "Mapping from JDBC Types to Java Object Types", in *JDBC API Tutorial and Reference, Second Edition: Universal Data Access for the Java 2 Platform*, 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*.

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.

## 4.5 Parameter mapping

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  $N$ ,  $0$  (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 must 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.

The book *JDBC API Tutorial and Reference, Second Edition: Universal Data Access for the Java 2 Platform* 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 JDBC's data type mapping tables are effectively extended. A row ( $ST$ ,  $JT$ ) is considered to be an additional row in Table 47.3, *Mapping from JDBC Types to Java Object Types*, and a row ( $JT$ ,  $ST$ ) is considered to be an additional row in Table 47.4, *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

**Replace 16th 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 11.49, "<SQL-invoked routine>", in ISO/IEC 9075-2, in Clause 16, "Embedded SQL", in ISO/IEC 9075-5, 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

**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 must 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 9.3, "<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.

#### 4.8 User-defined types

- 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.

Replace 8th paragraph A user-defined type is described by a user-defined type descriptor. A user-defined type descriptor includes:

- The name of the user-defined type (<user-defined type name>). This is the type designator of that type, used in type precedence lists (see Subclause 9.5, "Type precedence list determination", in ISO/IEC 9075-2).
- An indication of whether the user-defined type is ordered.
- The ordering form for the user-defined type (EQUALS, FULL, or NONE).
- The ordering category for the user-defined type (RELATIVE, COMPARABLE, MAP, or STATE).
- A <specific routine designator> identifying the ordering function, depending on the ordering category.
- If the user-defined type is a direct subtype of another user-defined type, then the name of that user-defined type.
- An indication of whether the user-defined type is instantiable or not instantiable.
- An indication of whether the user-defined type is final or not final.
- The transform descriptor of the user-defined type.
- If the user-defined type is a structured type, then:
  - Whether the referencing type of the structured type has a user-defined representation, a derived representation, or a system-defined representation.
  - If user-defined representation is specified, then the type descriptor of the representation type of the referencing type of the structured type; otherwise, if derived representation is specified, then the list of attributes.

NOTE 4 — "user-defined representation", "derived representation", and "system-defined representation" of a reference type are defined in Subclause 4.10, "Reference types", in ISO/IEC 9075-2.

- An indication of whether the user-defined type is an external Java data type.

Insert following the 8th 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>.

## 4.8 User-defined types

- 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 5 – 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>.

## 4.8.1 User-defined type comparison and assignment

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.

## 4.8.2 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.

### 4.8.3 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 2.0 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.

#### 4.8.3.1 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> must specify “implements `java.io.Serializable`” and must 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 must correspond to public fields of the corresponding Java class, which must 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).

#### 4.8.3.2 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 2.0. The Java class referenced in the <external Java class clause> of the <user-defined type definition> must specify “implements `java.sql.SQLData`” and must provide a no-arg constructor.



## 4.8 User-defined types

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.

### 4.8.3.3 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

Replace 1st paragraph A privilege authorizes a given category of <action> to be performed on a specified base table, view, column, domain, character set, collation, transliteration, user-defined type, trigger, SQL-invoked routine, or JAR by a specified <authorization identifier>.

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, 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, 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 6 – 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 must 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 must 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 must 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 must 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

### 5.1 <token> and <separator>

#### 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

### 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 7 – The rules of Java parsing and lexical analysis are specified in *The Java Language Specification, Second Edition*.
- 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 SR14) 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.

(Blank page)

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 6 Scalar expressions

### 6.1 <method invocation>

#### Function

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

#### Format

*No additional Format items.*

#### Syntax Rules

- 1) Insert after SR2) 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>

### 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

No additional General Rules.

### Conformance Rules

- 1) Insert after CR1) Without Feature J571, “NEW operator”, the schema identified by the implicit or explicit <schema name> of <routine name> *RN* immediately contained in <routine invocation> immediately contained in <new specification> shall not contain a user-defined type whose user-defined type name is *RN* that is an external Java data type. 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.

## 7 Predicates

### 7.1 <comparison predicate>

#### Function

Specify a comparison of two row values.

#### Format

*No additional Format items.*

#### Syntax Rules

NOTE 8 – Replace Note 126 The comparison form and comparison categories included in the user-defined type descriptors of both *UDT1* and *UDT2* are constrained to be the same — they must be the same throughout a type family. 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 *UDT1* and *UDT2* are constrained to be equivalent; if the comparison category is MAP, they are not constrained to be equivalent.

NOTE 9 – Replace Note 127 If the comparison form is FULL, then the comparison category is constrained to be COMPARABLE, RELATIVE, or MAP; if the comparison form is EQUALS, then the comparison category is also permitted to be STATE.

#### Access Rules

No additional Access Rules.

#### General Rules

- 1) Insert after GR 1)b)iii) and its subrules If the comparison category of *UDT<sub>x</sub>* is COMPARABLE, then:
  - a) The subject SQL data type must be an external Java data type. Let *JC* be the subject Java class of that external Java data type.  
 NOTE 10 – Syntax Rules in Subclause 9.10, “<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 Java `int`.
  - b) Let *XJV* be the value of *X* in the associated JVM. Let *YJV* be the value of *Y* in that associated JVM.
  - c)  $X = Y$

**7.1 <comparison predicate>**

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

### 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

None.

## 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 11 – 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 *JR*, 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>

### Function

Invoke an SQL-invoked routine.

### Format

No additional Format items.

### Syntax Rules

- 1) Insert before SR9) 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 8.5, “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 GR3a) 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 GR5) 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 12 – The subject Java method of an external Java routine is defined in Subclause 8.5, “Java routine signature determination”.
- 3) Replace the first paragraph of GR5) If  $R$  is an external routine that is not an external Java routine, then:
- 4) Replace the first paragraph of GR6)c)ii) If  $R$  is not a static field method, then:
- 5) Insert before GR9d) If  $R$  specifies PARAMETER STYLE JAVA, then for  $i$  ranging from 1 (one) to  $PN$ , let the effective SQL parameter list  $ESPL$  of  $R$  be the list of values  $CPV_i$  in order.
- 6) Replace the first paragraph of GR9f)iii)1) If  $R$  is not an external Java routine and  $R$  is not an array-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 GR9f)iii)2) If  $R$  is an external Java routine, 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 13 – The “subject static field” of an SQL-invoked method is defined in Subclause 8.5, “Java routine signature determination”.
  - 2) Let *ERT* be the effective returns data type of *R*.  
NOTE 14 – “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 15 – If *UIS* is *SERIALIZABLE*, then, as described in Subclause 9.3, “<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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.
        - 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 16 – If *UIS* is *SQLDATA*, then, as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

- 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 17 – 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<sub>i</sub>* be the *i*-th data type in *JPDL*.
  - 4) For *i* ranging from 1 (one) to *EN*, if *JP<sub>i</sub>* is of an array class, then let *JP<sub>i</sub>* be the component type of *JP<sub>i</sub>*.  
NOTE 18 – The component type of a Java array is defined in *The Java Language Specification, Second Edition*.
  - 5) For *i* ranging from 1 (one) to *EN*, if *ESP<sub>i</sub>* is the SQL null value and if *JP<sub>i</sub>* 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<sub>i</sub>* is a user-defined type, then let the most specific type of *ESP<sub>i</sub>* 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:
      - 1) If *UIS* is *SERIALIZABLE*, then:
        - 1) The subject Java class *SJCU*'s method `readObject( )` is executed to convert the value of *ESP<sub>i</sub>* to a Java object, the value of *PD<sub>i</sub>*.
        - 2) The method of execution of the subject Java class's implementation of `readObject( )` is implementation-defined.
- NOTE 19 – If *UIS* is *SERIALIZABLE*, then, as described in Subclause 9.3, "<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

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 20 – If  $UIS$  is SQLDATA, then, as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

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 *The Java Language Specification, Second Edition*.

NOTE 21 –  $PAD_i$  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 *The Java Language Specification, Second Edition*.

NOTE 22 –  $PAD_i$  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 tempU = JCLSN.JMN ( PD<sub>1</sub> , . . . , PD<sub>JPEN</sub> ) ;*

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

*SJCEN tempU = JCLSN.JMN ( ) ;*

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

- I) If *JPEN* is greater than 1 (one), then the following Java statement is effectively executed:

*JCLSN PD<sub>1</sub> = new JCLSN ( PD<sub>2</sub> , . . . , PD<sub>JPEN</sub> ) ;*

- II) If *JPEN* equals 1 (one), then the following Java statement is effectively executed:

*JCLSN PD<sub>1</sub> = new JCLSN ( ) ;*

- D) Otherwise:

- 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 1 (one), then the following Java statement is effectively executed:

*SJCEN tempU = PD<sub>1</sub> . JMN ( PD<sub>2</sub> , . . . , PD<sub>JPEN</sub> ) ;*

- IV) If *JPEN* equals 1 (one), then the following Java statement is effectively executed:

*SJCEN tempU = PD<sub>1</sub> . JMN ( ) ;*

**NOTE 23** – The Java method effectively executed by either the Java statement *SJCEN tempU = PD<sub>1</sub> . JMN ( PD<sub>2</sub> , . . . , PD<sub>JPEN</sub> ) ;* or the Java statement *SJCEN tempU = PD<sub>1</sub> . JMN ( ) ;* is determined based on the value of *PD<sub>1</sub>* according to Java's rules for overriding by instance methods, as specified in *The Java Language Specification, Second Edition*.

- 8) Insert before GR9)f)iii)4) If  $R$  is an external Java routine, then the scope and persistence of any modifications of class variables made before the completion of any execution of  $P$  is implementation-dependent.
- 9) Replace GR9)f)iii)7) If the language specifies ADA (respectively C, COBOL, FORTRAN, JAVA, MUMPS, PASCAL, PLI) and  $P$  is not a standard-conforming Ada program (respectively C, COBOL, Fortran, Java, MUMPS, Pascal, PL/I program), then the results of any execution of  $P$  are implementation-dependent.
- 10) Insert after GR9)f)iii)7) If  $R$  is an external Java routine and the execution of  $P$  completes with an uncaught Java exception  $E$ , then an exception condition is raised as specified in Subclause 15.1, “Class and subclass values for uncaught Java exceptions”, and no further General Rules of this Subclause are applied.
- 11) Replace the first paragraph of GR9)g)ii) If  $R$  is not an external Java routine, then for  $i$  varying from 1 (one) to  $EN$ , the value of  $ESP_i$  is set to the value of  $PD_i$ .
- 12) Insert after GR9)h)i)3) If  $R$  is an external Java routine that is not a type-preserving function, then let  $ERT$  be the effective returns data type of  $R$ . The returned value of  $P$ ,  $tempU$ , is processed as follows:
  - a) Case:
    - i) If  $ERT$  is a user-defined type, then:
      - 1) Let  $SJCE$  be the most specific Java class of the value of  $tempU$ , 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  $ERT$ .
      - 2) Let  $UIS$  be the <interface specification> specified by the user-defined type descriptor of  $STU$ .
      - 3) Case:
        - A) If  $UIS$  is SERIALIZABLE, then:
          - I) The subject Java class  $SJCE$ 's method `writeObject()` is executed to convert the Java value of  $tempU$  to the SQL value  $SSFV$  of user-defined type  $STU$ .
          - II) The method of execution of the subject Java class's implementation of `writeObject()` is implementation-defined.
        - NOTE 24 – If  $UIS$  is SERIALIZABLE, then, as described in Subclause 9.3, “<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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.
        - B) If  $UIS$  is SQLDATA, then:
          - I) The subject Java class  $SJCE$ 's method `writeSQL()` is executed to convert the Java value of  $tempU$  to the SQL value  $SSFV$  of user-defined type  $STU$ .
          - II) The method of execution of the subject Java class's implementation of `writeSQL()` is implementation-defined.
      - NOTE 25 – If  $UIS$  is SQLDATA, then as described in Subclause 9.3, “<user-defined type definition>”, the descriptor's subject Java class implements the Java interface

**ISO/IEC 9075-13:2002 (E)**  
**8.3 <routine invocation>**

`java.sql.SQLData` and defines that interface's `writeSQL( )` method as described by the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

ii) Otherwise, the value of *SSFV* is set to the value of *tempU*.

b) Let *RV* be *SSFV*.

13) Insert after GR9)h)i)3) If *R* is an external Java routine that is a type-preserving function, then let *ERT* be the effective returns data type of *R*. The returned value of *P*, *PD*<sub>1</sub>, is processed as follows:

- a) Let *SJCE* be the most specific Java class of the value of *PD*<sub>1</sub>, 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*.
- b) Let *UIS* be the <interface specification> specified by the user-defined type descriptor of *STU*.

Case:

i) If *UIS* is *SERIALIZABLE*, then:

- 1) The subject Java class *SJCE*'s method `writeObject( )` is executed to convert the Java value of *PD*<sub>1</sub> to the SQL value *SSFV* of user-defined type *STU*.
- 2) The method of execution of the subject Java class's implementation of `writeObject( )` is implementation-defined.

NOTE 26 – If *UIS* is *SERIALIZABLE*, then as described in Subclause 9.3, “<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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

ii) If *UIS* is *SQLDATA*, then:

- 1) The subject Java class *SJCE*'s method `writeSQL( )` is executed to convert the Java value of *PD*<sub>1</sub> to the SQL value *SSFV* of user-defined type *STU*.
- 2) The method of execution of the subject Java class's implementation of `writeSQL( )` is implementation-defined.

NOTE 27 – If *UIS* is *SQLDATA*, then as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

c) Let *RV* be *SSFV*.

14) Insert after GR9)i)ii) If *R* specifies *PARAMETER STYLE JAVA*, then each parameter that is either an output SQL parameter or both an input SQL parameter and an output SQL parameter is processed as follows:

- a) Let *P*<sub>*i*</sub> be the *i*-th SQL parameter of *R* and let *T*<sub>*i*</sub> be the declared type of *P*<sub>*i*</sub>.
- b) *EPV*<sub>*i*</sub> is set to the value of *PD*<sub>*i*</sub>[0].

Case:

i) If  $T_i$  is a user-defined type, then:

- 1) Let  $SJCE$  be the most specific Java class of the value of  $EPV_i$ , and let  $STU$  be the user-defined type whose subject Java class is  $SJCE$  and whose user-defined type is  $T_i$  or is a subclass of  $T_i$ .
- 2) Let  $UIS$  be the <interface specification> specified by the user-defined type descriptor of  $STU$ .

Case:

A) If  $UIS$  is `SERIALIZABLE`, then:

I) The subject Java class  $SJCE$ 's method `writeObject()` is executed to convert the Java value of  $EPV_i$  to the SQL value  $CPV_i$  of the user-defined type  $STU$ .

II) The method of execution of the subject Java class's implementation of `writeObject()` is implementation-defined.

NOTE 28 – If  $UIS$  is `SERIALIZABLE`, then as described in Subclause 9.3, “<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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

B) If  $UIS$  is `SQLDATA`, then:

I) The subject Java class  $SJCE$ 's method `writeSQL()` is executed to convert the Java value of  $EPV_i$  to the SQL value  $CPV_i$  of user-defined type  $STU$ .

II) The method of execution of the subject Java class's implementation of `writeSQL()` is implementation-defined.

NOTE 29 – If  $UIS$  is `SQLDATA`, then as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

ii) Otherwise,  $CPV_i$  is set to  $EPV_i$ .

- 15) Replace GR11)b) If  $R$  is not an external Java routine, then let  $OPN$  be the actual number of result set cursors declared in the body of the subject routine that remain open when control is returned to  $INV$ .
- 16) Insert after GR11)b) If  $R$  is an external Java routine, then let  $RSN$  be a set containing the first element of each of the  $JPEN-EN$  arrays generated above for result set mappable parameters, let  $RS$  be the elements of  $RSN$  that are not equal to the Java null value, and let  $OPN$  be the number of elements in  $RS$ .
- 17) Insert before GR11)d) If  $R$  is an external Java routine, then:
  - a) If the JDBC connection object that created any element of  $RS$  is closed, then the effect is implementation-defined.
  - b) If any element of  $RS$  is not an object returned by a connection to the current SQL system and SQL session, then the effect is implementation-defined.

### 8.3 <routine invocation>

- 18) Replace GR11)d) If  $R$  is not an external Java routine, then let  $FRC$  be the ordered set of result set cursors that remain open when  $PR$  returns to  $INV$ . Let  $FRC_i$ ,  $1 \text{ (one)} \leq i \leq RTN$ , be the  $i$ -th cursor in  $FRC$ , let  $FRCN_i$  be the <cursor name> that identifies  $FRC_i$ , and let  $RS_i$  be the result set of  $FRC_i$ .
- 19) Insert after GR11)d) If  $R$  is an external Java routine, then let  $FRC$  be a copy of the elements of  $RS$  that remain open in the order that they were opened in SQL. Let  $FRC_i$ ,  $1 \text{ (one)} \leq i \leq RTN$ , be the  $i$ -th cursor in  $FRC$ , let  $FRCN_i$  be the <cursor name> that identifies  $FRC_i$ , and let  $RCS_i$  be the result set of  $FRC_i$ .
- 20) Replace GR11)h) If  $R$  is not an external Java routine, then a completion condition is raised: *warning — dynamic result sets returned*.
- 21) Insert after GR11)h) If  $R$  is an external Java routine, then for each result set  $RS_i$  in  $RS$ , close  $RS_i$  and close the statement object that created  $RS_i$ .
- 22) Insert before GR12)h) If  $R$  is an external Java routine, then whether the call of  $P$  returns update counts as defined in JDBC is implementation-defined.

### Conformance Rules

- 1) Insert this CR Without Feature J611, “References”, conforming SQL language shall not contain a <reference expression> or a <right arrow>.

## 8.4 <language clause>

### Function

Specify a standard programming language.

### Format

```
<language name> ::=  
    !! All alternatives from ISO/IEC 9075-2  
    | JAVA
```

### Syntax Rules

No additional Syntax Rules.

### Access Rules

No additional Access Rules.

### General Rules

- 1) Replace GR 1) With the exception of the language JAVA, whose standard is specified in *The Java Language Specification, Second Edition*, the standard programming language specified by the <language clause> is defined in the International Standard identified by the <language name> keyword. Table 17, "Standard programming languages", in ISO/IEC 9075-2 specifies the relationship.

### Conformance Rules

No additional Conformance Rules.

## 8.5 Java routine signature determination

### Function

Specify rules for how a Java method's signature is determined if it is not explicitly specified, and how it is validated, based either on information specified when creating an external Java routine or external Java data type, or on contents of descriptors available when invoking an SQL routine.

### Syntax Rules

- 1) Let *CNTXT*, *i*, and *SR* respectively be the syntactic element, the method specification index, and the subject routine (if any) specified in an application of this Subclause.
- 2) Information needed by later rules of this Subclause is gathered based on the context in which this Subclause is executed, as follows.

Case:

- a) If *CNTXT* specifies <SQL-invoked routine>, then:
  - i) Let *JN*, *JCLSN*, *JMN*, and *JPDL* respectively be the <jar name>, <Java class name>, <Java method name>, and <Java parameter declaration list> contained in <external Java reference string>.
  - ii) Let *SPDL* be <SQL parameter declaration list>.
  - iii) If <SQL-invoked routine> contains <schema procedure>, then:
    - 1) If DYNAMIC RESULT SETS *N* is specified for some *N* greater than 0 (zero), then let *DRSN* be *N*.
    - 2) Otherwise let *DRSN* be 0 (zero).
- b) If *CNTXT* specifies <user-defined type definition>, then:
  - i) Let *UDTD* be the <user-defined type definition>, let *UDTB* be the <user-defined type body> immediately contained in *UDTD*, and let *UDTN* be the <schema-resolved user-defined type name> immediately contained in *UDTB*.
  - ii) Let *JN* and *JCLSN* respectively be the <jar name> and <Java class name> contained in <external Java type clause> contained in *UDTB*.
  - iii) For the purposes of parameter mapping as defined in Subclause 4.5, "Parameter mapping", the remaining rules in this Subclause are performed as if the descriptor for the user-defined type defined by *UDTD* was already available in the SQL-session. That descriptor describes the type as having the name *UDTN*, being an external Java data type, and having the <jar and class name> specified in *UDTD*.
  - iv) Let *MS<sub>i</sub>* be the *i*-th <method specification> in the <method specification list> contained by *UDTB*.
  - v) Let *SRT* be the SQL <data type> specified in the RETURNS clause of *MS<sub>i</sub>*.
  - vi) Let *DRSN* be 0 (zero).



- vii) If  $MS_i$  immediately contains <static field method spec>, then:
  - 1) Let  $QJFN$  be the <qualified Java field name> of  $MS_i$ .
  - 2) Let  $FI$  be the <Java identifier> contained in <Java field name> contained in  $QJFN$ .
  - 3) If  $QJFN$  specifies a <Java class name>, then let  $SFC$  be that class name; otherwise, let  $SFC$  be  $JCLSN$ .
  - 4) Let  $SPDL$  be the <SQL parameter declaration list>
    - ( )
- viii) If  $MS_i$  does not immediately contain <static field method spec>, then:
  - 1) Let  $JMN$  and  $JPDL$  respectively be the <Java method name> and <Java parameter declaration list> contained in <Java method and parameter declarations> contained in  $MS_i$ .
  - 2) Let  $SPDL$  be the augmented SQL parameter declaration list  $NPL_i$  of  $MS_i$ .
- c) Otherwise, descriptors are available.
  - i) Let  $SRD$  be the routine descriptor of  $SR$ .
  - ii) If  $SRD$  indicates that the SQL-invoked routine is an SQL-invoked method, then:
    - 1) Let  $SRUDT$  be the user-defined type whose descriptor contains  $SR$ 's corresponding method specification descriptor  $MSD$ , and let  $SRUDTD$  be the user-defined type descriptor of  $SRUDT$ .
    - 2) Let  $JN$  and  $JCLSN$  respectively be the <jar name> and <Java class name> contained by  $SRUDTD$ 's <jar and class name>.
    - 3) Let  $SRT$  be the SQL <returns data type> specified in  $MSD$ .
    - 4) Let  $DRSN$  be 0 (zero).
    - 5) If  $MSD$  indicates that it is a static field method, then:
      - A) Let  $FI$  be the <Java identifier> contained in the <Java field name> of  $MSD$ .
      - B) Let  $SFC$  be the <Java class name> of  $MSD$ .
      - C) Let  $SPDL$  be the <SQL parameter declaration list>
        - ( )
    - 6) If  $MSD$  indicates that it is not a static field method, then:
      - A) Let  $JMN$  and  $JPDL$  respectively be the Java method name composed of the package, class, and name of the Java routine contained in  $MSD$  and the Java parameter declaration list contained in the signature contained in  $MSD$ .
      - B) Let  $SPDL$  be the augmented SQL parameter declaration list of  $MSD$ .

## 8.5 Java routine signature determination

- iii) If *SRD* indicates that the SQL-invoked routine is an SQL-invoked procedure or an SQL-invoked regular function, then:
  - 1) Let *JN*, *JCLSN*, *JMN*, and *JPDL* respectively be the <jar name>, <Java class name>, <Java method name>, and <Java parameter declaration list> contained in <external Java reference string> contained in the <external routine name> of *SRD*.
  - 2) Let *SPDL* be a SQL parameter declaration list composed of the SQL-invoked routine's SQL parameters contained in *SRD*, specified with the descriptors list of the <SQL parameter name>, if specified, the <data type>, the ordinal position, and an indication of whether the SQL parameter is an input SQL parameter, an output SQL parameter, or both an input SQL parameter and an output SQL parameter.
  - 3) If the SQL-invoked routine is an SQL-invoked procedure, then let *DRSN* be the maximum number of dynamic result sets as indicated by *SRD*; otherwise, let *DRSN* be 0 (zero).
  - 4) If the SQL-invoked routine is an SQL-invoked regular function, then let *SRT* be the SQL <returns data type> specified in *MSD*; otherwise, let *SRT* be "void".

## 3) Case:

- a) If *JMN* is "main" and *CNTXT* does not specify <user-defined type definition> or contain <method invocation>, then:
  - i) If *CNTXT* specifies <SQL-invoked routine>, then it shall contain <schema procedure> and shall not contain <dynamic result set characteristic>.
  - ii) If *CNTXT* contains <routine invocation> then it shall contain <call statement>.
- iii) If a Java parameter declaration list *JPDL* is specified, then it shall be the following:
 

```
( java.lang.String[ ] )
```
- iv) If a Java parameter declaration list is not specified, then let *JPDL* be the following:
 

```
( java.lang.String[ ] )
```
- v) *SPDL* shall specify either:
  - 1) A single parameter that is an SQL ARRAY of CHARACTER or an ARRAY of CHARACTER VARYING. At runtime, this parameter is passed as a Java array of `java.lang.String`.  
NOTE 30 – This <SQL parameter declaration> can only be specified if the SQL system supports Feature S201, "SQL routines on arrays".
  - 2) Zero or more parameters, each of which is CHARACTER or CHARACTER VARYING. At runtime, these parameters are passed a Java array of `java.lang.String` (with possibly zero elements).
- vi) Let *JCS* be the set of visible Java methods of class *JCLSN* in JAR *JN* whose method names are "main" and whose Java parameter data types list is *JPDL*.  
NOTE 31 – "visible" is defined in Subclause 4.5, "Parameter mapping".

b) Otherwise:

i) Let *SPN* and *JPN* be, respectively, the number of <SQL parameter declaration>s in *SPDL* and the number of <Java data type>s in *JPDL*.

ii) If *JPDL* specifies a <Java parameter declaration list>, then:

1) If *i* is greater than 0 (zero) and *MS<sub>i</sub>* specifies INSTANCE or CONSTRUCTOR or if *SRD* indicates the SQL-invoked routine is an SQL-invoked method and *MSD* indicates it is an instance method or a constructor, then prefix the Java parameter declaration list *JPDL* with the necessary subject parameter as follows.

Case:

A) If *JPDL* contains one or more <Java data type>s, then prefix the list of <Java data type>s immediately contained in <Java parameters> immediately contained in *JPDL* with

*JCLSN* ,

B) Otherwise, replace *JPDL* with the <Java parameter declaration list>

( *JCLSN* )

2) For each <SQL parameter declaration> *SP* in *SPDL*, let *ST* be the <data type> of *SP* and let *JT* be the corresponding <Java data type> in *JPDL*.

A) If *SP* specifies IN, or does not specify an explicit <parameter mode>, then:

I) If *SP* is not an SQL array, then *JT* and *ST* shall be simply mappable or object mappable.

II) If *SP* is an SQL array, then *JT* and *ST* shall be array mappable.

B) If *SP* specifies OUT or INOUT, then *JT* and *ST* shall be output mappable.

NOTE 32 – “simply mappable”, “object mappable”, and “array mappable” are defined in Subclause 4.5, “Parameter mapping”.

3) Case:

A) If *DRSN* is greater than 0 (zero), then *JPN* shall be greater than *SPN*, and each <Java data type> in *JPDL* whose ordinal position is greater than *SPN* shall be result set mappable.

B) Otherwise *JPN* shall be equivalent to *SPN*.

iii) If a Java parameter declaration list is not specified, then determine the first *SPN* members of the Java parameter declaration list *JPDL* from *SPDL* as follows:

1) For each parameter *SP* of *SPDL* whose <parameter mode> is IN, or that does not specify an explicit <parameter mode>, if *SP* is not an SQL array, then let the corresponding Java parameter data type of *SP* be the corresponding Java data type of the <parameter type> of *SP*; if *SP* is an SQL array, then let *JT* be the corresponding Java data type of the <parameter type> of *SP*, and let the corresponding Java parameter data type of *SP* be an array of *JT*, that is, be *JT*[ ].

NOTE 33 – The “corresponding Java parameter data type” of *SP* is defined in Subclause 4.5, “Parameter mapping”.

## 8.5 Java routine signature determination

- 2) For each parameter *SP* of *SPDL* whose <parameter mode> is INOUT or OUT, let *JT* be the corresponding Java data type of the <parameter type> of *SP*, and let the corresponding Java parameter data type of *SP* be an array of *JT*, that is, be *JT*[ ].
- 3) The <Java parameters> of *JPDL* is a list of the corresponding Java parameter data types of *SPDL*.

NOTE 34 – *JPDL* does not specify parameter names. That is, the parameter names of the Java method do not have to match the SQL parameter names.

- iv) The subject Java field of <static field method spec>s or the set of candidate visible Java methods are determined as follows:

Case:

- 1) If *CNTXT* specifies <SQL-invoked routine> or if *SRD* indicates that the SQL-invoked routine is an SQL-invoked procedure or an SQL-invoked regular function, then:
  - A) If *DRSN* is greater than 0 (zero), then:
    - I) Let *SPN* and *JPN* be, respectively, the number of <SQL parameter declaration>s in *SPDL* and the number of <Java data type>s in *JPDL*.
    - II) If *SPN* is equivalent to *JPN*, then *JPDL* was originally not specified; let *JCS* be the set of visible Java methods of class *JCLSN* in JAR *JN* whose method names are *JMN*, whose first *SPN* parameter data types are those of *JPDL*, and whose last *K* parameter data types, for some positive *K*, are result set mappable.
    - III) If *SPN* is less than *JPN*, then *JPDL* was originally specified; let *JCS* be the set of visible Java methods of class *JCLSN* in JAR *JN* whose method names are *JMN*, whose Java parameter data types list is *JPDL*.
  - B) If *DRSN* is 0 (zero), then let *JCS* be the set of visible Java methods of class *JCLSN* in JAR *JN* whose method names are *JMN*, whose Java parameter data types list is *JPDL*.
- 2) If *CNTXT* specifies <user-defined type definition> or if *SRD* indicates that the SQL-invoked routine is an SQL-invoked method then:
  - A) If *i* is greater than 0 (zero) and *MS<sub>i</sub>* contains <static field method spec>, or if *MSD* indicates that it is a static field method, then:
    - I) *FI* shall be the name of a field of *SFC*. Let *JSF* be that field.
    - II) *JSF* shall be a public static field.
    - III) Let *JFT* be the Java data type of *JSF*.
    - IV) *SRT* and *JFT* shall be simply mappable or object mappable.  
NOTE 35 – “simply mappable” and “object mappable” are defined in Subclause 4.5, “Parameter mapping”.
    - V) *JSF* is the subject static field of the SQL-invoked method defined by *MS<sub>i</sub>*.  
NOTE 36 – The subject Java class may contain fields and methods (public and private) for which no corresponding attribute or method is specified.

## 8.5 Java routine signature determination

- B) If  $i$  is greater than 0 (zero) and  $MS_i$  does not immediately contain <static field method spec>, or if  $MSD$  indicates that it is not a static field method, then:

I) Case:

- 1) If  $i$  is greater than 0 (zero) and  $MS_i$  specifies INSTANCE or CONSTRUCTOR, or if  $MSD$  indicates it is an instance method or a constructor, then  $JPDL$  contains the augmented Java parameter declaration list for this method. Remove the subject parameter from the Java parameter declaration list  $JPDL$  to create the unaugmented Java parameter declaration list  $UAJPDL$ , as follows:

Case:

- a) If  $JPDL$  contains two or more <Java data type>s, then copy all  $JPDL$  to  $UAJPDL$ , omitting the first <Java data type>  $JCLSN$ , and its associated “,”.
- b) Otherwise, set  $UAJPDL$  to the <Java parameter declaration list>

( )

- 2) Otherwise copy  $JPDL$  to  $UAJPDL$ .

- II) Using Java overloading resolution, specified by *The Java Language Specification, Second Edition*, let  $JCS$  be the set of visible Java methods of class  $JCLSN$  in JAR  $JN$  or the supertypes of that class whose method names are  $JMN$  and whose Java parameter data types list is  $UAJPDL$ .

NOTE 37 – “visible” is defined in Subclause 4.5, “Parameter mapping”.

- III) If  $i$  is greater than 0 (zero) and  $MS_i$  specifies STATIC, or  $MSD$  indicates that STATIC was specified, then remove from  $JCS$  any Java method that is not static. Otherwise, remove from  $JCS$  any static Java method.
- IV) If  $i$  is greater than 0 (zero) and  $MS_i$  specifies CONSTRUCTOR, or  $MSD$  indicates that CONSTRUCTOR was specified, then remove from  $JCS$  any Java method that is not a constructor. Otherwise, remove from  $JCS$  any Java method that is a constructor.

- 4) The subject Java method is determined as follows:

Case:

- a) If  $CNTXT$  specifies <SQL-invoked routine> or if  $SRD$  indicates that the SQL-invoked routine is an SQL-invoked procedure or an SQL-invoked regular function, then:
- i)  $JCS$  shall contain exactly one Java method. Let  $JM$  be that Java method. The SQL-invoked routine is associated with  $JM$ .
- ii)  $JM$  is the subject Java method of the SQL-invoked routine.
- b) If  $CNTXT$  specifies <user-defined type definition> or if  $SRD$  indicates that the SQL-invoked routine is an SQL-invoked method then, if  $i$  is greater than 0 (zero) and  $MS_i$  does not

## 8.5 Java routine signature determination

immediately contain <static field method spec>, or if *MSD* indicates that it is not a static field method then:

i) *JCS* shall contain exactly one Java method. Let *JM* be that Java method. The <Java method name> is referred to as the *corresponding Java method name* of <method name>.

ii) *JM* is the *subject Java method* of the SQL-invoked method.

5) The result data type of the SQL-invoked routine is validated as follows:

Case:

a) If *CNTXT* specifies <SQL-invoked routine> or if *SRD* indicates that the SQL-invoked routine is an SQL-invoked procedure or an SQL-invoked regular function, then let *JRT* be the Java returns data type of *JM*.

i) If *JM* is an SQL-invoked procedure, then *JRT* shall be void.

ii) If *JM* is an SQL-invoked regular function, then *JRT* and *SRT* shall be simply mappable or object mappable.

b) If *CNTXT* specifies <user-defined type definition> or if *SRD* indicates that the SQL-invoked routine is an SQL-invoked method then, if *i* is greater than 0 (zero) and *MS<sub>i</sub>* does not immediately contain <static field method spec>, or if *MSD* indicates that it is not a static field method, then let *JRT* be the Java returns data type of *JM*. If SELF AS RESULT is not specified then *JRT* and *SRT* shall be simply mappable or object mappable.

NOTE 38 – “simply mappable” and “object mappable” are defined in Subclause 4.5, “Parameter mapping”.

c) Otherwise, let *JRT* be the Java data type of the subject static field. *JRT* and *SRT* shall be simply mappable or object mappable.

## 9 Schema definition and manipulation

### 9.1 <table definition>

#### Function

Define a persistent base table, a created local temporary table, or a global temporary table.

#### Format

*No additional Format items.*

#### Syntax Rules

- 1) Insert after SR 10)a) *ST* shall not be an external Java data type whose descriptor specifies an <interface specification> of `SERIALIZABLE`.

#### Access Rules

No additional Access Rules.

#### General Rules

No additional General Rules.

#### Conformance Rules

No additional Conformance Rules.



## 9.2 <view definition>

### Function

Define a viewed table.

### Format

*No additional Format items.*

### Syntax Rules

- 1) Insert after SR 21)c) *ST* shall not be an external Java data type whose descriptor specifies an <interface specification> of `SERIALIZABLE`.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

No additional Conformance Rules.

## 9.3 <user-defined type definition>

### Function

Define a user-defined type.

### Format

```
<user-defined type body> ::=
    <user-defined type name>
    [ <subtype clause> ]
    [ <external Java type clause> ]
    [ AS <representation> ]
    [ <instantiateable clause> ]
    <finality>
    [ <reference type specification> ]
    [ <ref cast option> ]
    [ <cast option> ]
    [ <method specification list> ]

<external Java type clause> ::=
    <external Java class clause>
    LANGUAGE JAVA
    <interface using clause>

<interface using clause> ::=
    [ USING <interface specification> ]

<interface specification> ::=
    SQLDATA
    | SERIALIZABLE

<method specification> ::=
    !! All alternatives from ISO/IEC 9075-2
    | <static field method spec>

<method characteristic> ::=
    !! All alternatives from ISO/IEC 9075-2
    | <external Java method clause>

<static field method spec> ::=
    STATIC METHOD <method name> <left paren> <right paren>
    <static method returns clause>
    [ SPECIFIC <specific method name> ]
    <external variable name clause>

<static method returns clause> ::=
    RETURNS <data type>

<external variable name clause> ::=
    EXTERNAL VARIABLE NAME <character string literal>

<external Java class clause> ::= EXTERNAL NAME <character string literal>

<external Java method clause> ::= EXTERNAL NAME <character string literal>

<Java method and parameter declarations> ::=
    <Java method name> [ <Java parameter declaration list> ]
```

## Syntax Rules

- 1) Insert after SR 3) If <external Java type clause> is specified, then *UDT* is an *external Java data type*.
- 2) Replace SR 6)k)i) The <supertype name> immediately contained in the <subtype clause> shall identify the descriptor of some structured type *SST*. *UDT* is a direct subtype of *SST*, and *SST* is a direct supertype of *UDT*. If *UDT* is an external Java data type, then *SST* shall be an external Java data type, and the subject Java class of *UDT* shall be a direct subclass of the subject Java class of *SST*. If *UDT* is not an external Java data type, then *SST* shall not be an external Java data type.
- 3) Insert before SR 7) If <external Java type clause> is specified, then:
  - a) Let *VJC* be the value of the <character string literal> immediately contained in <external Java class clause>; *VJC* shall conform to the Format and Syntax Rules of <jar and class name>. The Java class identified by <Java class name> in the JAR identified by <jar id> in their immediately containing <jar and class name> is *UDT*'s *subject Java class*.  
NOTE 39 – The subject Java class of *UDT* can be the subject Java class of other external Java data types. Each such external Java data type is distinct from other such data types.
  - b) *UDT*'s subject Java class shall be a public class and shall implement the Java interface `java.io.Serializable` or the Java interface `java.sql.SQLData` or both.
  - c) If an <interface using clause> is not explicitly specified, then an implementation-defined <interface specification> is implicit.
  - d) If `SERIALIZABLE` is specified, then the subject Java class shall implement the Java interface `java.io.Serializable`. The method `java.io.Serializable.writeObject( )` is effectively used to convert a Java object to an SQL representation, and the method `java.io.Serializable.readObject( )` is effectively used to convert an SQL representation to a Java object.
  - e) If `SQLDATA` is specified, then the subject Java class shall implement the Java interface `java.sql.SQLData` as defined in JDBC. The method `java.sql.SQLData.writeSQL( )` is effectively used to convert a Java object to an SQL representation, and the method `java.sql.SQLData.readSQL( )` is effectively used to convert an SQL representation to a Java object.
  - f) <overriding method specification> shall not be specified.
  - g) A <representation> that is a <predefined type> shall not be specified.
  - h) `SELF AS LOCATOR` shall not be specified.
  - i) <locator indication> shall not be specified.
- 4) Insert before SR 7) If <external Java type clause> is not specified, then:
  - a) <method specification> shall not specify <static field method spec>.
  - b) <method characteristic> shall not specify <external Java method clause>.
  - c) The <language clause> immediately contained in <method characteristic> shall not specify `JAVA`.

- 5) Insert after SR 7)a) If *UDT* is an external Java data type, then it is implementation-defined whether validation of the explicit or implicit <Java parameter declaration list> is performed by <user-defined type definition> or when the corresponding SQL-invoked method is invoked.
- 6) Insert after SR 7)b)iv.1)5) in the TC If *UDT* is an external Java data type, then the <Java identifier> immediately contained in <Java method name> of  $MS_i$  shall be equivalent to the <Java identifier> immediately contained in the <class identifier> immediately contained in <jar and class name> of *UDT*.
- 7) Insert after SR 7)b)x)4)B) *UDT* shall not be an external Java data type.
- 8) Insert after SR 7)b)xi)3) *UDT* shall not be an external Java data type.
- 9) Insert after SR 7)b)xiv) If  $MS_i$  specifies <static field method spec>, then:
  - a)  $MS_i$  specifies a *static field method*.
  - b) Let *VSF* be the value of the <character string literal> simply contained in <static field method spec>; *VSF* shall conform to the Format and Syntax Rules of <qualified Java field name>.  
 NOTE 40 – <static field method spec> defines a static method of the user-defined type that returns the value of the Java static field specified by the <qualified Java field name>. This is a shorthand that provides read-only SQL access to static fields of the subject Java class or a superclass of the subject Java class.
- 10) Replace SR 7)b)xv)1) The <method characteristics> of  $MS_i$  shall contain at most one <language clause>, at most one <parameter style clause>, at most one <deterministic characteristic>, at most one <SQL-data access indication>, and at most one <null-call clause>. If *UDT* is an external Java data type then, with the exception of the implicit <original method specification>s generated for the observer and mutator functions of each attribute, the <method characteristics> of  $MS_i$  shall not contain the <method characteristic>s <language clause> or <parameter style clause> and shall contain exactly one <external Java method clause>. For an external Java data type, both <language clause> and <parameter style clause> implicitly specify JAVA.
- 11) Insert after SR 7)b)xv)1) If *UDT* is an external Java data type, then let *VMP* be the value of the <character string literal> immediately contained in <external Java method clause>; *VMP* shall conform to the Format and Syntax Rules of <Java method and parameter declarations>.
- 12) Replace SR 7)b)xv)2) If *UDT* is not an external Java data type and <language clause> is not specified, then LANGUAGE SQL is implicit.
- 13) Replace SR 7)b)xv)6)B)l) If <parameter style> is not specified and *UDT* is not an external Java data type, then PARAMETER STYLE SQL is implicit.
- 14) Insert after SR 7)b)xvi) If *UDT* is an external Java data type and validation of the <Java parameter declaration list> has been implementation-defined to be performed by <user-defined type definition>, then the Syntax Rules of Subclause 8.5, “Java routine signature determination” are applied with <user-defined type definition>, method specification index *i*, and no subject routine.

## Access Rules

No additional Access Rules.

## General Rules

- 1) Replace GR 5)g)vii) The explicit or implicit <parameter style> if the <language name> is SQL or JAVA.

## Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> in a <descriptor file>.
- 2) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> that specifies LANGUAGE JAVA outside of a <descriptor file>.
- 3) Insert this CR Without Feature J591, “Overloading”, the <method name> of a <method specification> shall not be equivalent to the <method name> of any other <method specification> in the same <user-defined type definition>.
- 4) Insert this CR Without Feature J641, “Static fields”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <static field method spec>.
- 5) Insert this CR Without Feature J541, “SERIALIZABLE”, a conforming <user-defined type definition> shall not specify SERIALIZABLE.
- 6) Insert this CR Without Feature J551, “SQLDATA”, a conforming <user-defined type definition> shall not specify SQLDATA.
- 7) Insert this CR A conforming implementation of this part of ISO/IEC 9075 shall support at least one of Feature J541, “SERIALIZABLE”, and Feature J551, “SQLDATA”.
- 8) Insert this CR Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> that specifies LANGUAGE JAVA.

## 9.4 <attribute definition>

### Function

Define an attribute of a structured type.

### Format

```
<attribute definition> ::=  
    <attribute name>  
    <data type>  
    [ <reference scope check> ]  
    [ <attribute default> ]  
    [ <collate clause> ]  
    [ <external Java attribute clause> ]  
  
<external Java attribute clause> ::=  
    EXTERNAL NAME <character string literal>
```

### Syntax Rules

- 1) [Insert after SR 1) ] If the <attribute definition> is contained in a <user-defined type definition> that is not an external Java data type or is contained in an <alter type statement>, then <attribute definition> shall not specify an <external Java attribute clause>.
- 2) [Insert after SR 1) ] If the <attribute definition> is contained in a <user-defined type definition> that specifies an external Java data type whose <interface specification> is **SERIALIZABLE**, then <attribute definition> shall specify an <external Java attribute clause>.
- 3) [Insert after SR 1) ] If an <external Java attribute clause> is specified, then let *VFN* be the value of the <character string literal> immediately contained in <attribute definition>; *VFN* shall conform to the Format and Syntax Rules of <Java field name>. The <Java field name> value of *VFN* is referred to as the *corresponding Java field name* of the <attribute name>.
- 4) [Insert after SR 1) ] If <attribute definition> is contained in a <user-defined type definition> that specifies an external Java data type, then <reference scope check>, <attribute default>, and <collate clause> shall not be specified.
- 5) [Insert after SR 1) ] If <attribute definition> is contained in a <user-defined type definition> that specifies an external Java data type, and if the <data type> specified in the <attribute definition> is a structured type *ST*, then *ST* shall be an external Java data type.

### Access Rules

No additional Access Rules.

### General Rules

- 1) [Insert after GR3(g) ] If the <attribute definition> contains an <external Java attribute clause>, then the corresponding Java field name of the <attribute name>.
- 2) [Replace GR6) ] An SQL-invoked method *OF* is created whose signature and result data type are as given in the descriptor of the original method specification of the observer function of *A*. Let *V* be a value in *UDT*.

Case

- a) If  $V$  is the SQL null value, then the invocation  $V.AN()$  of  $OF$  returns the result of:  
`CAST (NULL AS  $DT$ )`
- b) If  $UDT$  is not an external Java data type whose descriptor's <interface specification> specifies **SERIALIZABLE**, then  $V.AN()$  returns the value of  $A$  in  $V$ .
- c) If  $UDT$  is an external Java data type whose descriptor's <interface specification> specifies **SERIALIZABLE**, then the `readObject()` method of the subject Java class  $SJCE$  of  $V$  is effectively used to obtain a Java object from the value of  $V$ , the Java field that corresponds to the attribute specified in <Java field name> contained by <attribute definition> is accessed. Let  $JV$  and  $JCLS$  be respectively that Java value and its most specific Java class.

Case:

- i) If  $DT$  is a user-defined type, then:
  - 1) Let  $STU$  be the user-defined type whose subject Java class is  $JCLS$  and whose user-defined type is  $DT$  or is a subclass of  $DT$ .
  - 2) Let  $UIS$  be the <interface specification> specified by the user-defined type descriptor of  $STU$ .
  - 3) Case:
    - A) If  $UIS$  is **SERIALIZABLE**, then:
      - I) The subject Java class  $JCLS$ 's `writeObject()` method is executed to convert the Java value  $JV$  to the SQL value  $SV$  of user-defined type  $STU$ .
      - II) The method of execution of the subject Java class's implementation of `writeObject()` is implementation-defined.

NOTE 41 – If  $UIS$  is **SERIALIZABLE**, then, as described in Subclause 9.3, “<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 *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.
    - B) If  $UIS$  is **SQLDATA**, then:
      - I) The subject Java class  $JCLS$ 's `writeSQL()` method is executed to convert the Java value  $JV$  to the SQL value  $SV$  of user-defined type  $STU$ .
      - II) The method of execution of the subject Java class's implementation of `writeSQL()` is implementation-defined.

NOTE 42 – If  $UIS$  is **SQLDATA**, then, as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.
    - C) Otherwise, the value of  $SV$  is set to the value of  $JV$ .
  - 4)  $V.AN()$  returns the value of  $SV$ .



- 3) Replace GR7) An SQL-invoked method *MF* is created whose signature and result data type are as given in the descriptor of the original method specification of the mutator function of *A*. Let *V* be a value in *UDT* and let *AV* be a value in *DT*.

Case:

- a) If *V* is the SQL null value, then the invocation *V*.*AN*(*AV*) of *MF* raises an exception condition: *data exception — null instance used in mutator function*.
- b) If *UDT* is not an external Java data type whose descriptor's <interface specification> specifies *SERIALIZABLE*, then the invocation *V*.*AN*(*AV*) returns *V2* such that *V2*.*AN*( ) = *AV* and for every other observer function *ANX* of *UDT*, *V2*.*ANX*( ) = *V*.*ANX*( ) .
- c) If *UDT* is an external Java data type whose descriptor's <interface specification> specifies *SERIALIZABLE*, then the *readObject*( ) method of the subject Java class *SJCE* of *V* is effectively used to obtain a Java object from the value of *V*. Let *MST*, *JCLS*, and *Jtemp* be respectively the most specific type of *AV*, the subject Java class of *MST*, and the Java object obtained from *readObject*( ) .

i) Case:

- 1) If *MST* is a user-defined type, then:

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

- B) Case:

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

- 1) The subject Java class *JCLS*'s *readObject*( ) method is executed to convert the value of *AV* to a Java object *JV*.
      - 2) The method of execution of the subject Java class's implementation of *readObject*( ) is implementation-defined.

NOTE 43 – If *UIS* is *SERIALIZABLE*, then, as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

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

- 1) The subject Java class *JCLS*'s *readSQL*( ) method is executed to convert the value of *AV* to a Java object *JV*.
      - 2) The method of execution of the subject Java class's implementation of *readSQL*( ) is implementation-defined.

NOTE 44 – If *UIS* is *SQLDATA*, then, as described in Subclause 9.3, “<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 the *Java 2 Platform, Standard Edition, v1.2.2, API Specification*.

- 2) Otherwise, the value of *JV* is set to the value of *AV*.

- ii) The Java field of *Jtemp* that corresponds to the attribute specified in <Java field name> contained by <attribute definition> is assigned the value *JV*.



- iii) The subject Java class *SJCE* of *V*'s `writeObject( )` method is effectively used to obtain a SQL value *V2* from the Java value *Jtemp*.
- iv) The invocation *V.AN(AV)* returns *V2*.

### **Conformance Rules**

No additional Conformance Rules.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 9.5 <alter type statement>

### Function

Change the definition of a user-defined type.

### Format

*No additional Format items.*

### Syntax Rules

- 1) Insert after SR 1) *D* shall not be an external Java data type.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

No additional Conformance Rules.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 9.6 <drop data type statement>

### Function

Destroy a user-defined type.

### Format

*No additional Format items.*

### Syntax Rules

No additional Syntax Rules.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop type statement> in a <descriptor file>.
- 2) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop type statement> that drops an external Java type outside of a <descriptor file>.
- 3) Insert this CR Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop data type statement> that drops an external Java type.

## 9.7 <SQL-invoked routine>

### Function

Define an SQL-invoked routine.

### Format

```
<parameter style> ::=  
    !! All alternatives from ISO/IEC 9075-2  
    | JAVA  
  
<external Java reference string> ::=  
    <jar and class name> <period> <Java method name>  
    [ <Java parameter declaration list> ]
```

### Syntax Rules

- 1) [Insert after SR 3] If <SQL-invoked routine> specifies LANGUAGE JAVA, then no <SQL parameter declaration> specified in <SQL-invoked function> shall specify RESULT.
- 2) [Insert after SR 3] If <SQL-invoked routine> specifies LANGUAGE JAVA, then neither the <returns clause> contained in <SQL-invoked function> nor any <SQL parameter declaration> contained in an <SQL-invoked function> or <SQL-invoked procedure> shall contain <locator indication>.
- 3) [Insert after SR 3] If <SQL-invoked routine> specifies LANGUAGE JAVA, then <transform group specification> shall not be specified.
- 4) [Insert after SR 3] The maximum value of <maximum dynamic result sets> is implementation-defined.
- 5) [Replace SR 4)a] Let *UDTN* be the <user-defined type name> immediately contained in <method specification designator>. Let *UDT* be the user-defined type identified by *UDTN*. *UDT* shall not be an external Java type.
- 6) [Replace SR 5)a] <routine characteristics> shall contain at most one <language clause>, at most one <parameter style clause>, at most one <specific name>, at most one <deterministic characteristic>, at most one <SQL-data access indication>, at most one <null-call clause>, and at most one <dynamic result sets characteristic>. If LANGUAGE JAVA is specified, then <parameter style clause> shall specify <parameter style> JAVA.
- 7) [Replace SR 5)i] An <SQL-invoked routine> that specifies or implies LANGUAGE SQL is called an *SQL routine*; an <SQL-invoked routine> that does not specify LANGUAGE SQL is called an *external routine*. An external routine that specifies LANGUAGE JAVA is called an *external Java routine*.
- 8) [Insert after SR 5)i] If *R* is an external Java routine, then the <external routine name> immediately contained in <external body reference> shall specify a <character string literal>. Let *V* be the value of that <character string literal>. *V* shall conform to the Format and Syntax Rules of an <external Java reference string>.

NOTE 45 – *R* is defined by ISO/IEC 9075-2 to be the SQL-invoked routine specified by <SQL-invoked routine>.

## 9.7 <SQL-invoked routine>

- 9) Insert after SR 5)i) If *R* is an external Java routine, then the <Java method name> is the name of one or more Java methods in the class specified by <Java class name> in the JAR specified by <jar name>. The combination of <Java class name> and <Java method name> represent a fully qualified Java class name and method name. The method name can reference a method of the class, or a method of a superclass of the class.
- 10) Replace the first paragraph of SR 5)t)iii) If *R* is not an external Java routine, then  
Case:
- 11) Replace the first paragraph of SR 19)e) If PARAMETER STYLE GENERAL or PARAMETER STYLE JAVA is specified, then let the *effective SQL parameter list* be a list of *PN* parameters such that, for *i* ranging from 1 (one) to *PN*, the *i*-th effective SQL parameter list entry is defined as follows.
- 12) Replace SR 19)g) If <language clause> does not specify JAVA, then every <data type> in an effective SQL parameter list entry shall specify a data type listed in the SQL data type column for which the corresponding row in the host data type column is not 'None'.
- 13) Insert before SR 20)  
NOTE 46 – The rules for parameter type correspondence when LANGUAGE JAVA is specified are given in Subclause 4.5, "Parameter mapping".
- 14) Insert before SR 20) If *R* is an external Java routine, then it is implementation-defined whether validation of the explicit or implicit <Java parameter declaration list> is performed by <SQL-invoked routine> or when its SQL-invoked routine is invoked.
- 15) Insert before SR 20) If *R* is an external Java routine, and validation of the <Java parameter declaration list> has been implementation-defined to be performed by <SQL-invoked routine>, then the Syntax Rules of Subclause 8.5, "Java routine signature determination", are applied with the <SQL-invoked routine>, a method specification index of 0 (zero), and no subject routine.

## Access Rules

- 1) Insert after AR 1) If *R* is an external Java routine, then the applicable privileges for *A* shall include USAGE privilege on the JAR referenced in the <external Java reference string>.  
NOTE 47 – The references to *R* and *A* are defined in the Syntax Rules of Subclause 11.49, "<SQL-invoked routine>", in ISO/IEC 9075-2.

## General Rules

- 1) Replace GR 3)m) If the SQL-invoked routine is an external routine, then the routine descriptor includes an indication of whether the parameter passing style is PARAMETER STYLE JAVA, PARAMETER STYLE SQL, or PARAMETER STYLE GENERAL.
- 2) Replace the introductory text of GR 6)a)i) If *R* is not an external Java routine and the <SQL data access indication> in the descriptor of *R* is MODIFIES SQL DATA, READS SQL DATA, or CONTAINS SQL, then:

## Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify an <SQL-invoked routine> in a <descriptor file>.
- 2) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> that specifies LANGUAGE JAVA outside of a <descriptor file>.
- 3) Insert this CR Without Feature J581, “Output parameters”, in a conforming implementation of this part of ISO/IEC 9075, <SQL-invoked routine> shall not specify <parameter mode> OUT or INOUT.
- 4) Insert this CR Without Feature J631, “Java signatures”, a <Java parameter declaration list> shall be equivalent to the default Java method signature as determined in Subclause 8.5, “Java routine signature determination”.
- 5) Insert this CR The SQL data types recognized by JDBC are a superset of those defined by ISO/IEC 9075-2. Without Feature J521, “JDBC data types”, a <Java data type> shall have a corresponding SQL data type.
- 6) Insert this CR Without Feature J621, “external Java routines”, conforming implementations of this part of ISO/IEC 9075 shall not specify an <SQL-invoked routine> that specifies LANGUAGE JAVA.

## 9.8 <alter routine statement>

### Function

Alter a characteristic of an SQL-invoked routine.

### Format

*No additional Format items.*

### Syntax Rules

- 1) Insert after SR 1) *SR* shall not be an external Java routine.

NOTE 48 – *SR* is defined to be the SQL-invoked routine identified by the <alter routine statement>.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

No additional Conformance Rules.

## 9.9 <drop routine statement>

### Function

Destroy an SQL-invoked routine.

### Format

*No additional Format items.*

### Syntax Rules

- 1) Insert after SR 7) If *SR* is an external Java routine and <drop routine statement> is contained in a <descriptor file>, then <drop routine statement> shall specify a <routine type> of PROCEDURE or of FUNCTION.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop routine statement> in a <descriptor file>.
- 2) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop routine statement> that drops an external Java routine outside of a <descriptor file>.
- 3) Insert this CR Without Feature J621, “external Java routines”, conforming implementations of this part of ISO/IEC 9075 shall not specify an <drop routine statement> that drops an external Java routine.



## 9.10 <user-defined ordering definition>

### Function

Define a user-defined ordering for a user-defined type.

### Format

```
<ordering category> ::=
    !! All alternatives from ISO/IEC 9075-2
    | <comparable category>

<comparable category> ::=
    RELATIVE WITH COMPARABLE INTERFACE
```

### Syntax Rules

- 1) Replace SR 4) If <comparable category>, <relative category>, or <state category> is specified, then *UDT* shall be a maximal supertype.
- 2) Insert before SR 6) If <comparable category> is specified, then *UDT* shall be an external Java data type. Let *JC* be the subject Java class of that external Java data type. *JC* shall implement the Java interface `java.lang.Comparable`.
- 3) Replace the introductory paragraph of SR 6)b) If <comparable category> is not specified, then:

### Access Rules

No additional Access Rules.

### General Rules

- 1) Replace GR 3) Case:
  - a) If <relative category> is specified, then the ordering category in the user-defined type descriptor of *UDT* is set to RELATIVE.
  - b) if <map category> is specified, then the ordering category in the user-defined type descriptor of *UDT* is set to MAP.
  - c) If <comparable category> is specified, then the ordering category in the user-defined type descriptor of *UDT* is set to COMPARABLE.
  - d) Otherwise, the ordering category in the user-defined type descriptor of *UDT* is set to STATE.

### Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined ordering definition> in a <descriptor file>.
- 2) Insert this CR Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined ordering definition> that defines an ordering for an external Java type.

**9.10 <user-defined ordering definition>**

- 3) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined ordering definition> that defines an ordering for an external Java type outside of a <deployment file>.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 9.11 <drop user-defined ordering statement>

### Function

Destroy a user-defined ordering method.

### Format

*No additional Format items.*

### Syntax Rules

No additional Syntax Rules.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop user-defined ordering statement> in a <descriptor file>.
- 2) Insert this CR Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop user-defined ordering statement> that drops an ordering for an external Java type.
- 3) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop user-defined ordering statement> that defines an ordering for an external Java type outside of a <deployment file>.

## 10 Access control

### 10.1 <grant privilege statement>

#### Function

Define privileges.

#### Format

*No additional Format Items.*

#### Syntax Rules

No additional Syntax Rules.

#### Access Rules

No additional Access Rules.

#### General Rules

No additional General Rules.

#### Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <grant privilege statement> in a <descriptor file>.
- 2) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <grant privilege statement> that grants USAGE privilege on a JAR outside of a <descriptor file>.

## 10.2 <privileges>

### Function

Specify privileges.

### Format

```
<object name> ::=  
    !! All alternatives from ISO/IEC 9075-2  
    | JAR <jar name>
```

### Syntax Rules

- 1) Replace SR 3) If <object name> specifies a <domain name>, <collation name>, <character set name>, <transliteration name>, <user-defined type name>, or <jar name>, then <privileges> shall specify USAGE. Otherwise, USAGE shall not be specified.

### Access Rules

No additional Access Rules.

### General Rules

No additional General Rules.

### Conformance Rules

- 1) Insert this CR Without Feature J561, “Jar privileges”, an <action> shall not specify USAGE on an <object name> that immediately contains a <jar name>.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 10.3 <revoke statement>

### Function

Destroy privileges and role authorizations.

### Format

No additional Format Items.

### Syntax Rules

- 1) [Replace SR 10)a)iii)4]  $P$  and  $D$  are both usage privilege descriptors. The action and the identified domain, character set, collation, transliteration, user-defined type, or JAR of  $P$  are the same as the action and the identified domain, character set, collation, transliteration, user-defined type, or JAR of  $D$ , respectively.
- 2) [Insert after SR 29)b)]  $DT$  is an external Java data type and the revoke destruction action would result in  $A1$  no longer having in its applicable privileges USAGE on the JAR whose <jar name> is contained in the <jar and class name> of the descriptor of  $DT$ .
- 3) [Insert after SR 34)] Let  $JR$  be any JAR descriptor included in  $S1$ .  $JR$  is said to be *impacted* if the revoke destruction action would result in  $A1$  no longer having in its applicable privileges USAGE privilege on a JAR whose name is contained in a <resolution jar> contained in the SQL-Java path of  $JR$ .
- 4) [Insert after SR 35)o)] If  $RD$  is an external Java routine, USAGE on the JAR whose <jar name> is contained in <external Java reference string> contained in the <external routine name> of the descriptor of  $RD$ .
- 5) [Replace SR 37)] If RESTRICT is specified, then there shall be no abandoned privilege descriptors, abandoned views, abandoned table constraints, abandoned assertions, abandoned domain constraints, lost domains, lost columns, lost schemas, impacted domains, impacted columns, impacted collations, impacted character sets, impacted JARs, abandoned user-defined types, forsaken column descriptors, forsaken domain descriptors, or abandoned routine descriptors.

### Access Rules

No additional Access Rules.

### General Rules

- 1) [Insert after GR 17)] If the object identified by <object name> of the <revoke statement> specifies <jar name>, let  $J$  be the JAR identified by that <jar name>. For every impacted JAR descriptor  $JR$  and for each <path element>  $PE$  contained in the SQL-Java path of  $JR$  whose immediately contained <resolution jar> is  $J$ , the SQL-Java path of the JAR descriptor  $JR$  is modified such that it does not contain  $PE$ .

### Conformance Rules

- 1) Insert this CR Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <revoke statement> in a <descriptor file>.
- 2) Insert this CR Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <revoke statement> that revokes USAGE privilege on a JAR outside of a <descriptor file>.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 11 Built-in procedures

### 11.1 SQLJ.INSTALL\_JAR procedure

#### Function

Install a set of Java classes into the current SQL catalog and schema.

#### Signature

```
SQLJ.INSTALL_JAR (
    url      IN    CHARACTER VARYING(L),
    jar      IN    CHARACTER VARYING(L),
    deploy   IN    INTEGER )
```

Where *L* is an implementation-defined integer value.

#### Access Rules

- 1) The privileges required to invoke the SQLJ.INSTALL\_JAR procedure are implementation-defined.

#### General Rules

- 1) The SQLJ.INSTALL\_JAR procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions. If an invocation of SQLJ.INSTALL\_JAR raises an exception condition, then the effect on the install actions is implementation-defined.
- 2) The values of the `url` parameter that are valid are implementation-defined, and may include URLs whose format is implementation-defined. If the value of the `url` parameter does not conform to implementation-defined restrictions and does not identify a valid JAR, then an exception condition is raised: *Java DDL — invalid URL*.
- 3) Let *J* be the value of the `jar` parameter. Let *TJ* be the value of
 

```
TRIM ( BOTH ' ' FROM J )
```

 If *TJ* does not conform to the Format and Syntax Rules of <jar name>, then an exception condition is raised: *Java DDL — invalid JAR name*.
- 4) Let *JN* be the explicitly or implicitly qualified <jar id> specified in the <jar name> specified by *TJ*.
- 5) If there is an installed JAR whose name is *JN*, then an exception condition is raised: *Java DDL — invalid JAR name*.
- 6) The JAR is installed and associated with the name *JN*. All contents of the JAR are installed, including both visible and non-visible Java classes, and other items contained in the JAR. This JAR becomes the *associated JAR* of each new class. The non-visible Java classes and other items can be referenced by other Java methods.
- 7) A privilege descriptor is created that defines the USAGE privilege on the JAR identified by the `jar` parameter to the <authorization identifier> that owns the schema identified by the implicit or explicit <schema name> of the `jar` parameter. The grantor for the privilege descriptor is set to the special grantor value “\_SYSTEM”. The privilege is grantable.



**11.1 SQLJ.INSTALL\_JAR procedure**

- 8) If the value of the `deploy` parameter is not zero, and if the JAR contains one or more deployment descriptor files, then the install actions implied by those instances are performed in the order in which the deployment descriptor files appear in the manifest.

NOTE 49 – Deployment descriptor files and their install actions are specified in Subclause 4.11.1, “Deployment descriptor files”.

**Conformance Rules**

- 1) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify non-zero values of the `deploy` parameter.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 11.2 SQLJ.REPLACE\_JAR procedure

### Function

Replace an installed JAR.

### Signature

```
SQLJ.REPLACE_JAR (
    url      IN  CHARACTER VARYING (L),
    jar      IN  CHARACTER VARYING (L) )
```

Where: *L* is an implementation-defined integer value.

### Access Rules

- 1) The privileges required to invoke the SQLJ.REPLACE\_JAR procedure are implementation-defined.
- 2) The current user shall be the owner of the JAR specified by the value of the `jar` parameter.

### General Rules

- 1) The SQLJ.REPLACE\_JAR procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions.
- 2) The values of the `url` parameter that are valid are implementation-defined, and may include URLs whose format is implementation-defined. If the value of `url` identifies a valid JAR, then refer to the classes in that JAR as the *new classes*. If the value of the `url` parameter does not identify a valid JAR, then an exception condition is raised: *Java DDL — invalid URL*.
- 3) Let *J* be the value of the `jar` parameter. Let *TJ* be the value of  

```
TRIM ( BOTH ' ' FROM J )
```

If *TJ* does not conform to the format of `<jar name>`, then an exception condition is raised: *Java DDL — invalid JAR name*.
- 4) Let *JN* be the explicitly or implicitly qualified `<jar id>` specified in the `<jar name>` specified by *TJ*.
- 5) If there is an installed JAR with `<jar name>` *JN*, then refer to that JAR as the *old JAR*. Refer to the classes in the old JAR as the *old classes*. If there is not an installed JAR with `<jar name>` *JN*, then an exception condition is raised: *Java DDL — attempt to replace uninstalled JAR*. Equivalence of JAR names is determined by the rules for equivalence of identifiers as specified in Subclause 5.2, "`<token>` and `<separator>`", in ISO/IEC 9075-2.
- 6) Let the *matching old classes* be the old classes whose fully qualified class names are the names of new classes and let the *matching new classes* be the new classes whose fully qualified class names are the names of old classes. Let the *unmatched old classes* be the old classes that are not matching old classes and let the *unmatched new classes* be the new classes that are not matching new classes.
- 7) Let the *dependent SQL routines* of a JAR be the routines whose descriptor's `<external routine name>` specifies an `<external Java reference string>` whose immediately contained `<jar name>` is equivalent to the JAR name of that JAR.

## 11.2 SQLJ.REPLACE\_JAR procedure

- 8) If any dependent SQL routine of the old JAR references a method in an unmatched old class, then an exception condition is raised: *Java DDL — invalid class deletion*.  
NOTE 50 – This rule prohibits deleting classes that are referenced by external Java routines. This prohibition does not, however, prevent deletion of classes that are referenced only indirectly by other Java classes.
- 9) For each dependent SQL routine of the old JAR that references a method in a matching old class, let *CS* be the <SQL-invoked routine> that created the SQL routine. If *CS* is not a valid <SQL-invoked routine> for the corresponding new routine, then an exception condition is raised: *Java DDL — invalid replacement*.
- 10) Let the *dependent SQL types* of a JAR file be the external Java data types that have as their subject Java class a Java class contained in that JAR.  
NOTE 51 – “subject Java class” is defined in Subclause 9.3, “<user-defined type definition>”.
- 11) If there are any dependent SQL types of the specified JAR file that are unmatched old classes, then an exception condition is raised: *Java DDL — invalid class deletion*.  
NOTE 52 – This rule prohibits deleting classes that are referenced by external Java data types. This prohibition does not, however, prevent deletion of classes that are referenced only indirectly by other Java classes.
- 12) For each dependent SQL type, let *CT* be the <user-defined type definition> that created the SQL type. If *CT* is not a valid <user-defined type definition> for the corresponding new class, then an exception condition is raised: *Java DDL — invalid replacement*.
- 13) The old JAR and all visible and non-visible old classes contained in it are deleted.
- 14) The new JAR and all visible and non-visible new classes are installed and associated with the specified <jar name>. That JAR becomes the *associated JAR* of each new class. All contents of the new JAR are installed, including both visible and non-visible Java classes, and other items contained in the JAR. The non-visible Java classes and other items can be referenced by other Java methods.
- 15) The effect of SQLJ.REPLACE\_JAR on currently executing SQL statements that use an SQL routine or structured type whose implementation has been replaced is implementation-dependent.

## Conformance Rules

None.

## 11.3 SQLJ.REMOVE\_JAR procedure

### Function

Remove an installed JAR and its classes.

### Signature

```
SQLJ.REMOVE_JAR (
    jar          IN    CHARACTER VARYING (L),
    undeploy     IN    INTEGER )
```

Where: *L* is an implementation-defined integer value.

### Access Rules

- 1) The privileges required to invoke the SQLJ.REMOVE\_JAR procedure are implementation-defined.
- 2) The current user shall be the owner of the JAR specified by the value of the *jar* parameter.

### General Rules

- 1) The SQLJ.REMOVE\_JAR procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions. If an invocation of SQLJ.REMOVE\_JAR raises an exception condition, then the effect on the remove actions is implementation-defined.
- 2) Let *J* be the value of the *jar* parameter. Let *TJ* be the value of  

```
TRIM ( BOTH ' ' FROM J )
```

If *TJ* does not conform to the format of <jar name>, then an exception condition is raised: *Java DDL — invalid JAR name*.
- 3) Let *JN* be the explicitly or implicitly qualified <jar id> specified in the <jar name> specified by *TJ*.
- 4) If there is an installed JAR with <jar name> *JN*, then refer to that JAR as the *old JAR*. Refer to the classes in the old JAR as the *old classes*. If there is not an installed JAR with <jar name> *JN*, then an exception condition is raised: *Java DDL — attempt to remove uninstalled JAR*. Equivalence of <jar name>*s* is determined by the rules for equivalence of identifiers as specified in Subclause 5.2, "<token> and <separator>", in ISO/IEC 9075-2.
- 5) If the value of the *undeploy* parameter is not 0 (zero), and if the JAR contains one or more deployment descriptor files, then the remove actions implied by those instances are performed in the reverse of the order in which the deployment descriptor files appear in the manifest.  
NOTE 53 — Deployment descriptor files and their remove actions are specified in Subclause 4.11.1, "Deployment descriptor files".  
NOTE 54 — These actions are performed prior to examining the condition specified in the following step.
- 6) Let the *dependent SQL routines* of a JAR be the routines whose descriptor's <external routine name> specifies an <external Java reference string> whose immediately contained <jar name> is equivalent to the name of that JAR.
- 7) If there are any dependent SQL routines of the specified JAR, then an exception condition is raised: *Java DDL — invalid class deletion*.

### 11.3 SQLJ.REMOVE\_JAR procedure

NOTE 55 – This rule prohibits deleting classes that are referenced by external Java routines. This prohibition does not, however, prevent deletion of classes that are referenced only indirectly by other Java classes.

- 8) Let the *dependent SQL types* of a JAR be the external Java data types that have as their subject Java class a Java class contained in that JAR.

NOTE 56 – “Subject Java class” is defined in Subclause 9.3, “<user-defined type definition>”.

- 9) If there are any dependent SQL types of the specified JAR, then an exception condition is raised: *Java DDL — invalid class deletion*.

NOTE 57 – This rule prohibits deleting classes that are referenced by external Java data types. This prohibition does not, however, prevent deletion of classes that are referenced only indirectly by other Java classes.

- 10) The specified JAR and all visible and non-visible classes contained in it are deleted.
- 11) The USAGE privilege on the specified JAR is revoked from all users that have it.
- 12) The effect of SQLJ.REMOVE\_JAR on currently executing SQL statements that use an SQL routine or structured type whose implementation has been removed is implementation-dependent.

### Conformance Rules

- 1) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify non-zero values of the `undeploy` parameter.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 11.4 SQLJ.ALTER\_JAVA\_PATH procedure

### Function

Alter the SQL-Java path of a JAR.

### Signature

```
SQLJ.ALTER_JAVA_PATH (  
    jar      IN      CHARACTER VARYING (L),  
    path     IN      CHARACTER VARYING (L) )
```

Where: *L* is an implementation-defined integer value.

### Access Rules

- 1) The privileges required to invoke the SQLJ.ALTER\_JAVA\_PATH procedure are implementation-defined.
- 2) The current user must be the owner of the JAR specified by the value of the *jar* parameter.
- 3) The current user must have the USAGE privilege on each JAR referenced in the *path* parameter.

### General Rules

- 1) The SQLJ.ALTER\_JAVA\_PATH procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions.
- 2) Let *J* be the value of the *jar* parameter. Let *TJ* be the value of  

```
TRIM ( BOTH ' ' FROM J )
```

If *TJ* does not conform to the format of <jar name>, then an exception condition is raised: *Java DDL — invalid JAR name*.
- 3) Let *JN* be the explicitly or implicitly qualified <jar id> specified in the <jar name> specified by *TJ*.
- 4) When the SQLJ.ALTER\_JAVA\_PATH procedure is called, the current catalog and schema at the time of the call are the default for each omitted <catalog name> and <schema name> in the <resolution jar>s of the *path* parameter. Those defaults apply to any subsequent use of the *path* parameter as specified below.
- 5) If the value of the *path* parameter does not conform to the format for <SQL Java path>, then an exception condition is raised: *Java DDL — invalid path name*.  
NOTE 58 – The *path* parameter can be an empty or all-blank string.
- 6) The value of the *path* parameter becomes the path associated with the JAR denoted by *JN*, replacing the current path (if any) associated with that JAR.
- 7) If an invocation of the SQLJ.ALTER\_JAVA\_PATH procedure raises an exception condition, then effect on the path associated with the JAR is implementation-defined.
- 8) The effect of SQLJ.ALTER\_JAVA\_PATH on SQL statements that have already been prepared or are currently executing is implementation-dependent.

### Conformance Rules

- 1) Without Feature J601, “SQL-Java paths”, conforming implementations of this part of ISO/IEC 9075 shall not contain invocations of the SQLJ.ALTER\_JAVA\_PATH procedure.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 12 Java topics

### 12.1 Java facilities supported by this part of ISO/IEC 9075

#### 12.1.1 Package `java.sql`

SQL systems that implement this part of ISO/IEC 9075 support the package `java.sql`, which is the JDBC driver, and all classes required by that package. The other Java packages supplied by SQL systems that implement this part of ISO/IEC 9075 are implementation-defined.

In an SQL system that implements this part of ISO/IEC 9075, the package `java.sql` supports the *default connection*. The default connection for a Java method invoked as an SQL routine has the following characteristics:

- The default connection is pre-allocated to provide efficient access to the database.
- The default connection is included in the current session and transaction.
- The authorization ID of the default connection is the current authorization ID.
- The JDBC AUTOCOMMIT setting of the default connection is *false*.

Other data source URLs that are supported by `java.sql` are implementation-defined.

#### 12.1.2 System properties

SQL systems that implement this part of ISO/IEC 9075 support the following system properties for use by the `getProperty` method of `java.lang.System`:

Table 2—System properties

Key	Description of associated value
<code>sqlj.defaultconnection</code>	If a Java method is executing in an SQL-implementation, then the String <code>"jdbc:default:connection"</code> <sup>1</sup>
<code>sqlj.runtime</code>	The class name of a runtime context class <sup>2</sup>
<sup>1</sup> Otherwise, the null value. <sup>2</sup> This class is a subclass of the class <code>sqlj.runtime.RuntimeContext</code> . The <code>getDefaultContext()</code> method of the class whose name is returned returns the default connection described in Subclause 12.1.1, "Package <code>java.sql</code> ".	



## 12.2 Deployment descriptor files

### Function

Supply information for actions to be taken by the `SQLJ.INSTALL_JAR` and `SQLJ.REMOVE_JAR` procedures.

### Model

A deployment descriptor file is a text file contained in a JAR, which is specified with the following property in the manifest for the JAR:

Name: *file\_name*  
SQLJDeploymentDescriptor: TRUE

### Properties

The text contained in a deployment descriptor file must have the following form:

```
<descriptor file> ::=
    SQLActions <left bracket> <right bracket> <equal sign>
    { [ <double quote> <action group> <double quote>
      [ <comma> <double quote> <action group> <double quote> ] ] }

<action group> ::=
    <install actions>
    | <remove actions>

<install actions> ::=
    BEGIN INSTALL [ <command> <semicolon> ]... END INSTALL

<remove actions> ::=
    BEGIN REMOVE [ <command> <semicolon> ]... END REMOVE

<command> ::=
    <SQL statement>
    | <implementor block>

<SQL statement> ::= !! See Description

<implementor block> ::=
    BEGIN <implementor name> <SQL token>... END <implementor name>

<implementor name> ::= <identifier>

<SQL token> ::= !! See Description
```

### Description

- 1) <descriptor file> shall contain at most one <install actions> and at most one <remove actions>.
- 2) The <command>s specified in the <install actions> (if any) and <remove actions> (if any) specify the *install actions* and *remove actions* of the deployment descriptor file, respectively.

- 3) An `<SQL statement>` specified in an `<install actions>` shall be either:
  - a) An `<SQL-invoked routine>` whose `<language clause>` specifies `JAVA`. The procedures and functions created by those statements are called the *deployed routines* of the deployment descriptor file.
  - b) A `<grant privilege statement>` that specifies the `EXECUTE` privilege for a deployed routine.
  - c) A `<user-defined type definition>` that specifies an `<external Java type clause>`. The types created by those statements are called the *deployed types* of the deployment descriptor file.
  - d) A `<grant privilege statement>` that specifies the `USAGE` privilege for a deployed type.
  - e) A `<user-defined ordering definition>` that specifies ordering for a deployed type.
- 4) When a deployment descriptor file is executed by a call of the `SQLJ.INSTALL_JAR` procedure, if the `<jar name>` of any `<external routine name>` or an `<SQL-invoked routine>` in an `<install actions>` is the `<jar name>` `"thisjar"`, then `"thisjar"` is effectively replaced with the `jar` parameter of the `SQLJ.INSTALL_JAR` procedure for purposes of that execution.
- 5) An `<SQL statement>` specified in a `<remove actions>` shall be either:
  - a) A `<drop routine statement>` for a deployed routine.
  - b) A `<revoke statement>` for the `EXECUTE` privilege on a deployed routine.
  - c) A `<drop data type statement>` for a deployed type.
  - d) A `<revoke statement>` for the `USAGE` privilege on a deployed type.
  - e) A `<drop user-defined ordering statement>` that specifies ordering for a deployed type.
- 6) An `<implementor block>` specifies implementation-specific install actions (remove actions) when specified in an `<install actions>` (`<remove actions>`).
- 7) An `<SQL token>` is an SQL lexical unit specified by the term `"<token>"` in Subclause 5.2, `"<token>` and `<separator>"`, in ISO/IEC 9075-2. That is, the comments, quotes, and double-quotes in an `<implementor block>` follow SQL token conventions.
- 8) An `<implementor name>` is an implementation-defined SQL identifier. The `<implementor name>`s specified following the `BEGIN` and `END` keywords shall be equivalent.
- 9) Whether an `<implementor block>` with a given `<implementor name>` contained in an `<install actions>` (`<remove actions>`) is interpreted as an install action (remove action) is implementation-defined. That is, an implementation may or may not perform install or remove actions specified by some other implementation.

NOTE 59 – The deployment descriptor file format corresponds to the more general notion of a properties file supporting indexed properties. Therefore, the deployment descriptor file can be used by the SQL-implementation to instantiate a Java Bean having an indexed property, `SQLActions`. You can then customize the resulting Java Bean instance with ordinary Java Bean tools. For example, you can change the SQL procedures or permissions by changing the routine descriptors stored in the `SQLActions` property. The SQL system can then use the customized Java Bean instance to generate a modified version of the deployment descriptor file to use in a revised version of the JAR.

(Blank page)

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 13 Information Schema

### 13.1 JAR\_JAR\_USAGE view

#### Function

Identify each JAR owned by a given user or role on which JARs defined in this catalog are dependent.

#### Definition

```
CREATE VIEW JAR_JAR_USAGE AS
  SELECT SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME,
         JAR_CATALOG, JAR_SCHEMA, JAR_NAME
  FROM ( DEFINITION_SCHEMA.JAR_JAR_USAGE JJU
        JOIN
          DEFINITION_SCHEMA.JARS J
        USING ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME ) )
  JOIN
    DEFINITION_SCHEMA.SCHEMATA S
  ON ( ( JJU.PATH_JAR_CATALOG, JJU.PATH_JAR_SCHEMA )
      = ( S.CATALOG_NAME, S.SCHEMA_NAME ) )
  WHERE ( SCHEMA_OWNER = CURRENT_USER
        OR
          SCHEMA_OWNER IN
            ( SELECT ROLE_NAME
              FROM ENABLED_ROLES ) )
  AND
    JAR_CATALOG =
      ( SELECT CATALOG_NAME
        FROM INFORMATION_SCHEMA.CATALOG_NAME ) ;

GRANT SELECT ON TABLE JAR_JAR_USAGE
  TO PUBLIC WITH GRANT OPTION;
```

#### Conformance Rules

- 1) Without Feature J652, "SQL/JRT Usage tables", conforming SQL language shall not reference INFORMATION\_SCHEMA.JAR\_JAR\_USAGE.

## 13.2 JARS view

### Function

Identify the installed JARs defined in this catalog that are accessible to the current user.

### Definition

```
CREATE VIEW JARS AS
  SELECT JAR_CATALOG, JAR_SCHEMA, JAR_NAME
  FROM DEFINITION_SCHEMA.JARS
 WHERE ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME, 'JAR' ) IN
        ( SELECT OBJECT_CATALOG, OBJECT_SCHEMA, OBJECT_NAME, OBJECT_TYPE
          FROM DEFINITION_SCHEMA.USAGE_PRIVILEGES
          WHERE GRANTEE IN
                ( 'PUBLIC', CURRENT_USER )
            OR
              GRANTEE IN
                ( SELECT ROLE_NAME
                  FROM ENABLED_ROLES ) )
 AND
        JAR_CATALOG =
        ( SELECT CATALOG_NAME
          FROM INFORMATION_SCHEMA.CATALOG_NAME ) ;

GRANT SELECT ON TABLE JARS
  TO PUBLIC WITH GRANT OPTION;
```

### Conformance Rules

- 1) Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA.JARS.

## 13.3 METHOD\_SPECIFICATIONS view

### Function

Identify the methods in the catalog that are accessible to a given user or role.

### Definition

Add columns EXTERNAL_NAME and IS_FIELD in ISO/IEC 9075-2
--

 Add “, EXTERNAL\_NAME, IS\_FIELD”  
to the end of the outermost <select list> of the <view definition>.

### Conformance Rules

- 1) 

Insert this CR
----------------

 Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA . METHOD\_SPECIFICATIONS . EXTERNAL\_NAME or INFORMATION\_SCHEMA . METHOD\_SPECIFICATIONS . IS\_FIELD.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 13.4 ROUTINE\_JAR\_USAGE view

### Function

Identify the JARs owned by a given user or role on which external Java routines defined in this catalog are dependent.

### Definition

```
CREATE VIEW ROUTINE_JAR_USAGE AS
  SELECT SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME,
         JAR_CATALOG, JAR_SCHEMA, JAR_NAME
  FROM ( DEFINITION_SCHEMA.ROUTINE_JAR_USAGE
        JOIN
          DEFINITION_SCHEMA.ROUTINES
        USING ( SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME ) )
  JOIN
    DEFINITION_SCHEMA.SCHEMATA S
  ON ( ( JAR_CATALOG, JAR_SCHEMA ) =
      ( S.CATALOG_NAME, S.SCHEMA_NAME ) )
  WHERE ( SCHEMA_OWNER = CURRENT_USER
        OR
          SCHEMA_OWNER IN
            ( SELECT ROLE_NAME
              FROM ENABLED_ROLES ) )
  AND
    SPECIFIC_CATALOG =
      ( SELECT CATALOG_NAME
        FROM INFORMATION_SCHEMA.CATALOG_NAME );

GRANT SELECT ON TABLE ROUTINE_JAR_USAGE
  TO PUBLIC WITH GRANT OPTION;
```

### Conformance Rules

- 1) Without Feature J652, "SQL/JRT Usage tables", conforming SQL language shall not reference INFORMATION\_SCHEMA.ROUTINE\_JAR\_USAGE.

## 13.5 TYPE\_JAR\_USAGE view

### Function

Identify the JARs owned by a given user or role on which external Java types defined in this catalog are dependent.

### Definition

```
CREATE VIEW TYPE_JAR_USAGE AS
  SELECT USER_DEFINED_TYPE_CATALOG AS UDT_CATALOG,
         USER_DEFINED_TYPE_SCHEMA AS UDT_SCHEMA,
         USER_DEFINED_TYPE_NAME AS UDT_NAME,
         JAR_CATALOG, JAR_SCHEMA, JAR_NAME
  FROM ( DEFINITION_SCHEMA.TYPE_JAR_USAGE
        JOIN
          DEFINITION_SCHEMA.TYPES
        USING ( USER_DEFINED_TYPE_CATALOG, USER_DEFINED_TYPE_SCHEMA,
              USER_DEFINED_TYPE_NAME ) )
  JOIN
    DEFINITION_SCHEMA.SCHEMATA S
  ON ( ( JAR_CATALOG, JAR_SCHEMA ) =
      ( S.CATALOG_NAME, S.SCHEMA_NAME ) )
 WHERE ( SCHEMA_OWNER = CURRENT_USER
        OR
          SCHEMA_OWNER IN
            ( SELECT ROLE_NAME
              FROM ENABLED_ROLES ) )
  AND
    USER_DEFINED_TYPE_CATALOG =
      ( SELECT CATALOG_NAME
        FROM INFORMATION_SCHEMA.CATALOG_NAME ) ;

GRANT SELECT ON TABLE TYPE_JAR_USAGE
  TO PUBLIC WITH GRANT OPTION;
```

### Conformance Rules

- 1) Without Feature J652, "SQL/JRT Usage tables", conforming SQL language shall not reference INFORMATION\_SCHEMA.TYPE\_JAR\_USAGE.



## 13.6 USER\_DEFINED\_TYPES view

### Function

Identify the user-defined types defined in this catalog that are accessible to a given user or role.

### Definition

Add columns EXTERNAL_NAME, EXTERNAL_LANGUAGE, and JAVA_INTERFACE in ISO/IEC 9075-2
--

Add “, EXTERNAL\_NAME, EXTERNAL\_LANGUAGE, JAVA\_INTERFACE” to the end of the outermost <select list> of the <view definition>.

### Conformance Rules

- 1) 

Insert this CR
----------------

 Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_NAME, INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_LANGUAGE, or INFORMATION\_SCHEMA . UDT\_S . JAVA\_INTERFACE.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 13.7 Short name views

### Function

Provide alternative views that use only identifiers that do not require Feature F391, “Long identifiers”.

### Definition

Replace view METHOD\_SPECS in ISO/IEC 9075-2

```
CREATE VIEW METHOD_SPECS
( SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME,
  UDT_CATALOG, UDT_SCHEMA, UDT_NAME,
  METHOD_NAME, IS_STATIC, IS_OVERRIDING,
  IS_CONSTRUCTOR, DATA_TYPE, CHAR_MAX_LENGTH,
  CHAR_OCTET_LENGTH, CHAR_SET_CATALOG, CHAR_SET_SCHEMA,
  CHARACTER_SET_NAME, COLLATION_CATALOG, COLLATION_SCHEMA,
  COLLATION_NAME, NUMERIC_PRECISION, NUMERIC_PREC_RADIX,
  NUMERIC_SCALE, DATETIME_PRECISION, INTERVAL_TYPE,
  INTERVAL_PRECISION, RETURN_UDT_CATALOG, RETURN_UDT_SCHEMA,
  RETURN_UDT_NAME, SCOPE_CATALOG, SCOPE_SCHEMA,
  SCOPE_NAME, MAX_CARDINALITY, DTD_IDENTIFIER,
  METHOD_LANGUAGE, PARAMETER_STYLE, IS_DETERMINISTIC,
  SQL_DATA_ACCESS, IS_NULL_CALL, TO_SQL_SPEC_CAT,
  TO_SQL_SPEC_SCHEMA, TO_SQL_SPEC_NAME, AS_LOCATOR,
  EXTERNAL_NAME, IS_FIELD, CREATED,
  LAST_ALTERED ) AS
SELECT SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME,
  USER_DEFINED_TYPE_CATALOG, USER_DEFINED_TYPE_SCHEMA, USER_DEFINED_TYPE_NAME,
  METHOD_NAME, IS_STATIC, IS_OVERRIDING,
  IS_CONSTRUCTOR, DATA_TYPE, CHARACTER_MAXIMUM_LENGTH,
  CHARACTER_OCTET_LENGTH, CHARACTER_SET_CATALOG, CHARACTER_SET_SCHEMA,
  CHARACTER_SET_NAME, COLLATION_CATALOG, COLLATION_SCHEMA,
  COLLATION_NAME, NUMERIC_PRECISION, NUMERIC_PRECISION_RADIX,
  NUMERIC_SCALE, DATETIME_PRECISION, INTERVAL_TYPE,
  INTERVAL_PRECISION, RETURN_UDT_CATALOG, RETURN_UDT_SCHEMA,
  RETURN_UDT_NAME, SCOPE_CATALOG, SCOPE_SCHEMA,
  SCOPE_NAME, MAXIMUM_CARDINALITY, DTD_IDENTIFIER,
  METHOD_LANGUAGE, PARAMETER_STYLE, IS_DETERMINISTIC,
  SQL_DATA_ACCESS, IS_NULL_CALL, TO_SQL_SPECIFIC_CATALOG,
  TO_SQL_SPECIFIC_SCHEMA, TO_SQL_SPECIFIC_NAME, AS_LOCATOR,
  EXTERNAL_NAME, IS_FIELD, CREATED,
  LAST_ALTERED
FROM INFORMATION_SCHEMA.METHOD_SPECIFICATIONS;
```

Replace view UDT\_S in ISO/IEC 9075-2

## 13.7 Short name views

```

CREATE VIEW UDT_S
( UDT_CATALOG,          UDT_SCHEMA,          UDT_NAME,
  UDT_CATEGORY,         IS_INSTANTIABLE,     IS_FINAL,
  ORDERING_FORM,        ORDERING_CATEGORY,   ORDERING_ROUT_CAT,
  ORDERING_ROUT_SCH,    ORDERING_ROUT_NAME,  REFERENCE_TYPE,
  DATA_TYPE,           CHAR_MAX_LENGTH,     CHAR_OCTET_LENGTH,
  CHAR_SET_CATALOG,     CHAR_SET_SCHEMA,     CHARACTER_SET_NAME,
  COLLATION_CATALOG,    COLLATION_SCHEMA,     COLLATION_NAME,
  NUMERIC_PRECISION,    NUMERIC_PREC_RADIX,   NUMERIC_SCALE,
  DATETIME_PRECISION,   INTERVAL_TYPE,        INTERVAL_PRECISION,
  SOURCE_DTD_ID,        REF_DTD_IDENTIFIER,  EXTERNAL_NAME,
  EXTERNAL_LANGUAGE,    JAVA_INTERFACE ) AS
SELECT USER_DEFINED_TYPE_CATALOG, USER_DEFINED_TYPE_SCHEMA, USER_DEFINED_TYPE_NAME,
  CATEGORY, IS_INSTANTIABLE, IS_FINAL,
  ORDERING_FORM, ORDERING_CATEGORY, ORDERING_ROUTINE_CATALOG,
  ORDERING_ROUTINE_SCHEMA, ORDERING_ROUTINE_NAME, REFERENCE_TYPE,
  DATA_TYPE, CHARACTER_MAXIMUM_LENGTH, CHARACTER_OCTET_LENGTH,
  CHARACTER_SET_CATALOG, CHARACTER_SET_SCHEMA, CHARACTER_SET_NAME,
  COLLATION_CATALOG, COLLATION_SCHEMA, COLLATION_NAME,
  NUMERIC_PRECISION, NUMERIC_PRECISION_RADIX, NUMERIC_SCALE,
  DATETIME_PRECISION, INTERVAL_TYPE, INTERVAL_PRECISION,
  SOURCE_DTD_IDENTIFIER, REF_DTD_IDENTIFIER, EXTERNAL_NAME,
  EXTERNAL_LANGUAGE, JAVA_INTERFACE
FROM INFORMATION_SCHEMA.USER_DEFINED_TYPES;

```

## Conformance Rules

- 1) Insert this CR Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA . METHOD\_SPECS . EXTERNAL\_NAME, INFORMATION\_SCHEMA . METHOD\_SPECS . IS\_FIELD, INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_NAME, INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_LANGUAGE, or INFORMATION\_SCHEMA . UDT\_S . JAVA\_INTERFACE.

## 14 Definition Schema

### 14.1 JAR\_JAR\_USAGE base table

#### Function

The JAR\_JAR\_USAGE table has one row for each JAR included in the SQL-Java path of a JAR.

#### Definition

```
CREATE TABLE JAR_JAR_USAGE (
  JAR_CATALOG      INFORMATION_SCHEMA.SQL_IDENTIFIER,
  JAR_SCHEMA        INFORMATION_SCHEMA.SQL_IDENTIFIER,
  JAR_NAME          INFORMATION_SCHEMA.SQL_IDENTIFIER,
  PATH_JAR_CATALOG  INFORMATION_SCHEMA.SQL_IDENTIFIER,
  PATH_JAR_SCHEMA   INFORMATION_SCHEMA.SQL_IDENTIFIER,
  PATH_JAR_NAME     INFORMATION_SCHEMA.SQL_IDENTIFIER,

  CONSTRAINT JAR_JAR_USAGE_PRIMARY_KEY
    PRIMARY KEY ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME,
                  PATH_JAR_CATALOG, PATH_JAR_SCHEMA, PATH_JAR_NAME ),

  CONSTRAINT JAR_JAR_USAGE_CHECK_REFERENCES_JARS
    CHECK ( PATH_JAR_CATALOG NOT IN
            ( SELECT CATALOG_NAME
              FROM SCHEMATA )
      OR
            ( PATH_JAR_CATALOG, PATH_JAR_SCHEMA, PATH_JAR_NAME ) IN
            ( SELECT JAR_CATALOG, JAR_SCHEMA, JAR_NAME
              FROM JARS ) ),

  CONSTRAINT JAR_JAR_USAGE_FOREIGN_KEY_JARS
    FOREIGN KEY ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME )
      REFERENCES JARS
)
```

#### Description

- 1) The JAR\_JAR\_USAGE table has one row for each JAR *JP* identified by a <jar name> contained in an <SQL Java path> associated with JAR *J*.
- 2) The values of JAR\_CATALOG, JAR\_SCHEMA, and JAR\_NAME are the <catalog name>, <unqualified schema name>, and <jar id>, respectively, of the <jar name> of the JAR (*J*) being described.
- 3) The values of PATH\_JAR\_CATALOG, PATH\_JAR\_SCHEMA, and PATH\_JAR\_NAME are the <catalog name>, <unqualified schema name>, and <jar id>, respectively, of the <jar name> of a JAR (*JP*) that is in the <SQL Java path> associated with JAR *J*.

## 14.2 JARS base table

### Function

The JARS table has one row for each installed JAR.

### Definition

```
CREATE TABLE JARS (
  JAR_CATALOG      INFORMATION_SCHEMA.SQL_IDENTIFIER,
  JAR_SCHEMA        INFORMATION_SCHEMA.SQL_IDENTIFIER,
  JAR_NAME          INFORMATION_SCHEMA.SQL_IDENTIFIER,

  CONSTRAINT JARS_PRIMARY_KEY
    PRIMARY KEY ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME ),

  CONSTRAINT JAR_FOREIGN_KEY_SCHEMATA
    FOREIGN KEY ( JAR_CATALOG, JAR_SCHEMA )
      REFERENCES SCHEMATA
)
```

### Description

- 1) The values of JAR\_CATALOG, JAR\_SCHEMA, and JAR\_NAME are the <catalog name>, <unqualified schema name>, and <jar id> of the <jar name> of the JAR being described.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 14.3 METHOD\_SPECIFICATIONS base table

### Function

The METHOD\_SPECIFICATIONS base table has one row for each method specification.

### Definition

Replace CONSTRAINT METHOD\_SPECIFICATIONS\_LANGUAGE\_CHECK in ISO/IEC 9075-2

```
CONSTRAINT METHOD_SPECIFICATIONS_LANGUAGE_CHECK
CHECK ( METHOD_LANGUAGE IN
( 'SQL', 'ADA', 'C', 'COBOL',
'FORTRAN', 'MUMPS', 'PASCAL', 'PLI', 'JAVA' ) )
```

Add two columns and two constraints in ISO/IEC 9075-2 Add the following <table element>s to the <table element list> of the METHOD\_SPECIFICATIONS base table:

```
EXTERNAL_NAME      INFORMATION_SCHEMA.CHARACTER_DATA,
IS_FIELD           INFORMATION_SCHEMA.CHARACTER_DATA

CONSTRAINT METHOD_SPECIFICATIONS_IS_FIELD_CHECK
CHECK ( IS_FIELD IN ( 'YES', 'NO' ) ),

CONSTRAINT METHOD_SPECIFICATIONS_METHOD_COMBINATIONS
CHECK ( ( ( METHOD_LANGUAGE = 'JAVA' )
AND
( EXTERNAL_NAME, IS_FIELD ) IS NOT NULL ) )
OR
( ( METHOD_LANGUAGE <> 'JAVA' )
AND
( ( EXTERNAL_NAME, IS_FIELD )
IS NULL ) ) ),

CONSTRAINT METHOD_SPECIFICATIONS_FIELD_COMBINATIONS
CHECK ( IS_FIELD = 'NO' OR IS_STATIC = 'YES' )
```

### Description

1)  Case:

- a) If the method being described is an external Java routine, then the value of EXTERNAL\_NAME is the <Java method and parameter declarations> specified in the <external Java method clause> for that external Java data type.
- b) If the method being described is a static field of an external Java type, then the value of EXTERNAL\_NAME is the <qualified Java field name> specified in the <static field method spec> of the method.
- c) Otherwise, the value of EXTERNAL\_NAME is the null value.

2)  Case:

- a) If the method being described is a static field of an external Java type, then the value of IS\_FIELD is 'YES'.
- b) If the method being described is an external Java type, then the value of IS\_FIELD is 'NO'.
- c) Otherwise, the value of IS\_FIELD is the null value.

## 14.4 ROUTINE\_JAR\_USAGE base table

### Function

The ROUTINE\_JAR\_USAGE table has one row for each external Java routine that names a JAR in an <external Java reference string>.

### Definition

```
CREATE TABLE ROUTINE_JAR_USAGE (
    SPECIFIC_CATALOG      INFORMATION_SCHEMA.SQL_IDENTIFIER,
    SPECIFIC_SCHEMA       INFORMATION_SCHEMA.SQL_IDENTIFIER,
    SPECIFIC_NAME         INFORMATION_SCHEMA.SQL_IDENTIFIER,
    JAR_CATALOG           INFORMATION_SCHEMA.SQL_IDENTIFIER,
    JAR_SCHEMA            INFORMATION_SCHEMA.SQL_IDENTIFIER,
    JAR_NAME              INFORMATION_SCHEMA.SQL_IDENTIFIER,

    CONSTRAINT ROUTINE_JAR_USAGE_PRIMARY_KEY
        PRIMARY KEY ( SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME,
                     JAR_CATALOG, JAR_SCHEMA, JAR_NAME ),

    CONSTRAINT JAR_JAR_USAGE_CHECK_REFERENCES_JARS
        CHECK ( JAR_CATALOG NOT IN
              ( SELECT CATALOG_NAME
                FROM Schemata )
        OR
              ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME ) IN
              ( SELECT JAR_CATALOG, JAR_SCHEMA, JAR_NAME
                FROM JARS ) ),

    CONSTRAINT JAR_JAR_USAGE_FOREIGN_KEY_ROUTINES
        FOREIGN KEY (SPECIFIC_CATALOG, SPECIFIC_SCHEMA, SPECIFIC_NAME )
        REFERENCES ROUTINES
)
```

### Description

- 1) The ROUTINE\_JAR\_USAGE table has one row for each external Java routine that names a JAR in an <external Java reference string>.
- 2) The values of SPECIFIC\_CATALOG, SPECIFIC\_SCHEMA, and SPECIFIC\_NAME are the <catalog name>, <unqualified schema name>, and <qualified identifier>, respectively, of the <specific name> of the external Java routine being described.
- 3) The values of JAR\_CATALOG, JAR\_SCHEMA, and JAR\_NAME are the <catalog name>, <unqualified schema name>, and <jar id>, respectively, of the <jar name> of the JAR being referenced in the external Java routine's <external Java reference string>.

## 14.5 ROUTINES base table

### Function

The ROUTINES table has one row for each SQL-invoked routine.

### Definition

Replace CONSTRAINT EXTERNAL_LANGUAGE_CHECK in ISO/IEC 9075-2
--

 Add “, 'JAVA'” to the <in value list> of valid EXTERNAL\_LANGUAGE values.

### Description

No additional Descriptions.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002



## 14.6 TYPE\_JAR\_USAGE base table

### Function

The TYPE\_JAR\_USAGE table has one row for each external Java type.

### Definition

```
CREATE TABLE TYPE_JAR_USAGE (
    USER_DEFINED_TYPE_CATALOG      INFORMATION_SCHEMA.SQL_IDENTIFIER,
    USER_DEFINED_TYPE_SCHEMA        INFORMATION_SCHEMA.SQL_IDENTIFIER,
    USER_DEFINED_TYPE_NAME          INFORMATION_SCHEMA.SQL_IDENTIFIER,
    JAR_CATALOG                     INFORMATION_SCHEMA.SQL_IDENTIFIER,
    JAR_SCHEMA                      INFORMATION_SCHEMA.SQL_IDENTIFIER,
    JAR_NAME                        INFORMATION_SCHEMA.SQL_IDENTIFIER,

    CONSTRAINT TYPE_JAR_USAGE_PRIMARY_KEY
        PRIMARY KEY (USER_DEFINED_TYPE_CATALOG, USER_DEFINED_TYPE_SCHEMA,
                     USER_DEFINED_TYPE_NAME, JAR_CATALOG, JAR_SCHEMA, JAR_NAME),

    CONSTRAINT TYPE_JAR_USAGE_CHECK_REFERENCES_JARS
        CHECK ( JAR_CATALOG NOT IN
                ( SELECT CATALOG_NAME
                  FROM SCHEMATA )
        OR
        ( JAR_CATALOG, JAR_SCHEMA, JAR_NAME ) IN
        ( SELECT JAR_CATALOG, JAR_SCHEMA, JAR_NAME
          FROM JARS ) ),

    CONSTRAINT TYPE_JAR_USAGE_FOREIGN_KEY_USER_DEFINED_TYPES
        FOREIGN KEY (USER_DEFINED_TYPE_CATALOG, USER_DEFINED_TYPE_SCHEMA,
                     USER_DEFINED_TYPE_NAME ) REFERENCES USER_DEFINED_TYPES
)
```

### Description

- 1) The TYPE\_JAR\_USAGE table has one row for each external Java type.
- 2) The values of USER\_DEFINED\_TYPE\_CATALOG, USER\_DEFINED\_TYPE\_SCHEMA, and USER\_DEFINED\_TYPE\_NAME are the <catalog name>, <unqualified schema name>, and <qualified identifier>, respectively, of the <user-defined type name> of the external Java type being described.
- 3) The values of JAR\_CATALOG, JAR\_SCHEMA, and JAR\_NAME are the <catalog name>, <unqualified schema name>, and <jar id>, respectively, of the <jar name> of the JAR being referenced in the external Java type's <external Java class clause>.

## 14.7 USAGE\_PRIVILEGES base table

### Function

The USAGE\_PRIVILEGES table has one row for each usage privilege descriptor. It effectively contains a representation of the usage privilege descriptors.

### Definition

Replace CONSTRAINT USAGE\_PRIVILEGES\_OBJECT\_TYPE\_CHECK in ISO/IEC 9075-2 Add “, ‘JAR’” to the <in value list> of valid OBJECT\_TYPE values.

Replace CONSTRAINT USAGE\_PRIVILEGES\_CHECK\_REFERENCES\_OBJECT in ISO/IEC 9075-2 Add the following to the end of the <query expression> contained in the <in predicate>:

```
UNION
  SELECT JAR_CATALOG, JAR_SCHEMA, JAR_NAME, 'JAR'
FROM JARS
```

### Description

- 1) 

Augment Description 4)
------------------------

JAR	The object to which the privilege applies is a JAR.
-----	---

## 14.8 USER\_DEFINED\_TYPES base table

### Function

The USER\_DEFINED\_TYPES table has one row for each user-defined type.

### Definition

Add three columns and one constraint

Add the following <table element>s to the <table element list> of the USER\_DEFINED\_TYPES base table:

```
EXTERNAL_NAME          INFORMATION_SCHEMA.CHARACTER_DATA,
EXTERNAL_LANGUAGE      INFORMATION_SCHEMA.CHARACTER_DATA
CONSTRAINT USER_DEFINED_TYPE_EXTERNAL_LANGUAGE_CHECK
CHECK ( EXTERNAL_LANGUAGE IN ( 'JAVA' ) ),
JAVA_INTERFACE         INFORMATION_SCHEMA.CHARACTER_DATA
CONSTRAINT USER_DEFINED_TYPE_JAVA_INTERFACE_CHECK
CHECK ( JAVA_INTERFACE IN ( 'SERIALIZABLE', 'SQLDATA' ) ),
CONSTRAINT USER_DEFINED_TYPES_COMBINATIONS
CHECK ( ( ( EXTERNAL_LANGUAGE = 'JAVA' ) AND
          ( EXTERNAL_NAME, JAVA_INTERFACE ) IS NOT NULL ) )
OR
      ( ( EXTERNAL_LANGUAGE, EXTERNAL_NAME, JAVA_INTERFACE )
        IS NULL ) )
```

Augment CONSTRAINT USER\_DEFINED\_TYPES\_ORDERING\_CATEGORY\_CHECK in ISO/IEC 9075-2

Add “, 'COMPARABLE'” to the <in value list> of valid ORDERING\_CATEGORY values.

### Description

1) Augment Description 7)

COMPARABLE                      Two values of this type may be compared with java.lang.Comparable's compareTo( ) method.

2) Insert this Description Case:

- a) If the user-defined type being described is an external Java data type, then the value of EXTERNAL\_NAME is the <jar and class name> specified in the <external Java class clause> for that external Java data type.
- b) Otherwise, the value of EXTERNAL\_NAME is the null value.

3) Insert this Description Case:

- a) If the user-defined type being described is an external Java data type, then the value of EXTERNAL\_LANGUAGE is 'JAVA'.
- b) Otherwise, the value of EXTERNAL\_LANGUAGE is the null value.

4) Insert this Description Case:

- a) If the user-defined type being described is an external Java data type, then the value of JAVA\_INTERFACE is the <interface specification> specified for that external Java data type.

- b) Otherwise, the value of JAVA\_INTERFACE is the null value.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

(Blank page)

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## 15 Status codes

### 15.1 Class and subclass values for uncaught Java exceptions

When the execution of a Java method completes with an uncaught Java exception  $E$ , then:

- 1) Let  $EM$  be the result of the Java method call `E.getMessage()`.
- 2)  $EM$  is the message text associated with the SQL exception.
- 3) Case:
  - a) If the class of  $E$  is `java.sql.SQLException`, then let  $SS$  be the result of the Java method call `E.getSQLState()`:
    - i) If the length of  $SS$  is 5 or more, and the first two characters of  $SS$  are '38' (corresponding to *external routine exception*), and the third, fourth, and fifth characters are not '000', then let  $C$  be '38' (corresponding to *external routine exception*) and let  $SC$  be the third, fourth, and fifth characters of  $SS$ .
    - ii) Otherwise, let  $C$  be '39' (corresponding to *external routine invocation exception*) and  $SC$  be '001' (corresponding to *invalid SQLSTATE returned*).
  - b) If the class of  $E$  is not `java.sql.SQLException`, then let  $C$  be '38' (corresponding to *external routine exception*) and  $SC$  be '000' (corresponding to *no subclass*).
- 4)  $C$  and  $SC$  are the class and subclass of the SQLSTATE for the SQL exception.

## 15.2 SQLSTATE

Table 3—SQLSTATE class and subclass values

Category	Condition	Class	Subcondition	Subclass
X	<i>All alternatives from ISO/IEC 9075-2</i> Java DDL <sup>1</sup>	46	<i>(no subclass)</i> invalid URL invalid JAR name invalid class deletion invalid replacement attempt to replace uninstalled JAR attempt to remove uninstalled JAR	000 001 002 003 005 00A 00B
X	Java execution <sup>1</sup>	46	<i>(no subclass)</i> invalid JAR name in path unresolved class name	000 102 103

<sup>1</sup>The Condition names “Java DDL” and “Java execution” are given the same Class code given to Condition name “OLB-specific error” in ISO/IEC 9075-10; there is no conflict with Subcondition values for the Class code.

## 16 Conformance

### 16.1 Claims of conformance

Insert this paragraph An implementation of this standard is *conformant* if it implements all capabilities specified in this standard that are not specified as being optional, and if it identifies which of those capabilities specified as being optional that it also implements.

Insert this paragraph In addition to the requirements of Subclause 8.2.3, "Claims of conformance", in ISO/IEC 9075-1, a claim of conformance to this part of ISO/IEC 9075 shall state:

1) Which of the following may be specified for <language clause> in an <SQL-invoked routine>:

- ADA
- C
- COBOL
- FORTRAN
- MUMPS
- PASCAL
- PLI
- SQL
- JAVA

At least one of these shall be specified.

2) Which of the following features is supported:

- Feature J621, "external Java routines"
- Feature J622, "external Java types"

At least one of these shall be specified.

3) Whether or not Feature J561, "JAR privileges", is supported.

4) For each of the following features that the implementation supports:

- Feature J621, "external Java routines"
- Feature J622, "external Java types"
- Feature J561, "JAR privileges"

specify in which of the following manners the feature is supported:

- Feature J511, "Commands"



**16.1 Claims of conformance**

- Feature J531, “Deployment”

For each of Feature J621, “external Java routines”, Feature J622, “external Java types”, and Feature J561, “JAR privileges”, at least one of Feature J511, “Commands” and Feature J531, “Deployment” shall be specified.

- 5) Which of the following features is supported:

- Feature J541, “SERIALIZABLE”
- Feature J551, “SQLDATA”

At least one of these shall be specified.

- 6) Which, if any, of the following features are supported:

- Feature J521, “JDBC data types”
- Feature J571, “NEW operator”
- Feature J581, “Output parameters”
- Feature J591, “Overloading”
- Feature J601, “SQL-Java paths”
- Feature J611, “References”
- Feature J621, “external Java routines”
- Feature J622, “external Java types”
- Feature J631, “Java signatures”
- Feature J641, “Static fields”
- Feature J651, “SQL/JRT Information Schema”
- Feature J652, “SQL/JRT Usage tables”

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## Annex A (informative)

### SQL Conformance Summary

Replicated paragraph The contents of this Annex summarizes all Conformance Rules, ordered by Feature ID and by Subclause.

- 1) Specifications for Feature J511, “Commands”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> that specifies LANGUAGE JAVA outside of a <descriptor file>.
  - b) Subclause 9.6, “<drop data type statement>”:
    - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop type statement> that drops an external Java type outside of a <descriptor file>.
  - c) Subclause 9.7, “<SQL-invoked routine>”:
    - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> that specifies LANGUAGE JAVA outside of a <descriptor file>.
  - d) Subclause 9.9, “<drop routine statement>”:
    - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop routine statement> that drops an external Java routine outside of a <descriptor file>.
  - e) Subclause 9.10, “<user-defined ordering definition>”:
    - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined ordering definition> that defines an ordering for an external Java type outside of a <deployment file>.
  - f) Subclause 9.11, “<drop user-defined ordering statement>”:
    - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop user-defined ordering statement> that defines an ordering for an external Java type outside of a <deployment file>.

- g) Subclause 10.1, “<grant privilege statement>”:
  - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <grant privilege statement> that grants USAGE privilege on a JAR outside of a <descriptor file>.
- h) Subclause 10.3, “<revoke statement>”:
  - i) Without Feature J511, “Commands”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <revoke statement> that revokes USAGE privilege on a JAR outside of a <descriptor file>.
- 2) Specifications for Feature J521, “JDBC data types”:
  - a) Subclause 9.7, “<SQL-invoked routine>”:
    - i) The SQL data types recognized by JDBC are a superset of those defined by ISO/IEC 9075-2. Without Feature J521, “JDBC data types”, a <Java data type> shall have a corresponding SQL data type.
- 3) Specifications for Feature J531, “Deployment”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> in a <descriptor file>.
  - b) Subclause 9.6, “<drop data type statement>”:
    - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop type statement> in a <descriptor file>.
  - c) Subclause 9.7, “<SQL-invoked routine>”:
    - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify an <SQL-invoked routine> in a <descriptor file>.
  - d) Subclause 9.9, “<drop routine statement>”:
    - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop routine statement> in a <descriptor file>.
  - e) Subclause 9.10, “<user-defined ordering definition>”:
    - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined ordering definition> in a <descriptor file>.
  - f) Subclause 9.11, “<drop user-defined ordering statement>”:
    - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop user-defined ordering statement> in a <descriptor file>.

- g) Subclause 10.1, “<grant privilege statement>”:
  - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <grant privilege statement> in a <descriptor file>.
- h) Subclause 10.3, “<revoke statement>”:
  - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <revoke statement> in a <descriptor file>.
- i) Subclause 11.1, “SQLJ.INSTALL\_JAR procedure”:
  - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify non-zero values of the `deploy` parameter.
- j) Subclause 11.3, “SQLJ.REMOVE\_JAR procedure”:
  - i) Without Feature J531, “Deployment”, conforming implementations of this part of ISO/IEC 9075 shall not specify non-zero values of the `undeploy` parameter.
- 4) Specifications for Feature J541, “SERIALIZABLE”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J541, “SERIALIZABLE”, a conforming <user-defined type definition> shall not specify `SERIALIZABLE`.
  - b) Subclause 9.3, “<user-defined type definition>”:
    - i) A conforming implementation of this part of ISO/IEC 9075 shall support at least one of Feature J541, “SERIALIZABLE”, and Feature J551, “SQLDATA”.
- 5) Specifications for Feature J551, “SQLDATA”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J551, “SQLDATA”, a conforming <user-defined type definition> shall not specify `SQLDATA`.
- 6) Specifications for Feature J561, “Jar privileges”:
  - a) Subclause 10.2, “<privileges>”:
    - i) Without Feature J561, “Jar privileges”, an <action> shall not specify `USAGE` on an <object name> that immediately contains a <jar name>.
- 7) Specifications for Feature J571, “NEW operator”:
  - a) Subclause 6.2, “<new specification>”:
    - i) Without Feature J571, “NEW operator”, the schema identified by the implicit or explicit <schema name> of <routine name> *RN* immediately contained in <routine invocation> immediately contained in <new specification> shall not contain a user-defined type whose user-defined type name is *RN* that is an external Java data type. 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.

- 8) Specifications for Feature J581, “Output parameters”:
  - a) Subclause 9.7, “<SQL-invoked routine>”:
    - i) Without Feature J581, “Output parameters”, in a conforming implementation of this part of ISO/IEC 9075, <SQL-invoked routine> shall not specify <parameter mode> OUT or INOUT.
- 9) Specifications for Feature J591, “Overloading”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J591, “Overloading”, the <method name> of a <method specification> shall not be equivalent to the <method name> of any other <method specification> in the same <user-defined type definition>.
- 10) Specifications for Feature J601, “SQL-Java paths”:
  - a) Subclause 8.2, “<SQL Java path>”:
    - i) Without Feature J601, “SQL-Java paths”, conforming SQL language shall not contain an <SQL Java path>.
  - b) Subclause 11.4, “SQLJ.ALTER\_JAVA\_PATH procedure”:
    - i) Without Feature J601, “SQL-Java paths”, conforming implementations of this part of ISO/IEC 9075 shall not contain invocations of the SQLJ.ALTER\_JAVA\_PATH procedure.
- 11) Specifications for Feature J611, “References”:
  - a) Subclause 8.3, “<routine invocation>”:
    - i) Without Feature J611, “References”, conforming SQL language shall not contain a <reference expression> or a <right arrow>.
- 12) Specifications for Feature J621, “external Java routines”:
  - a) Subclause 9.7, “<SQL-invoked routine>”:
    - i) Without Feature J621, “external Java routines”, conforming implementations of this part of ISO/IEC 9075 shall not specify an <SQL-invoked routine> that specifies LANGUAGE JAVA.
  - b) Subclause 9.9, “<drop routine statement>”:
    - i) Without Feature J621, “external Java routines”, conforming implementations of this part of ISO/IEC 9075 shall not specify an <drop routine statement> that drops an external Java routine.
- 13) Specifications for Feature J622, “external Java types”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined type definition> that specifies LANGUAGE JAVA.

- b) Subclause 9.6, “<drop data type statement>”:
  - i) Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop data type statement> that drops an external Java type.
- c) Subclause 9.10, “<user-defined ordering definition>”:
  - i) Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <user-defined ordering definition> that defines an ordering for an external Java type.
- d) Subclause 9.11, “<drop user-defined ordering statement>”:
  - i) Without Feature J622, “external Java types”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <drop user-defined ordering statement> that drops an ordering for an external Java type.
- 14) Specifications for Feature J631, “Java signatures”:
  - a) Subclause 9.7, “<SQL-invoked routine>”:
    - i) Without Feature J631, “Java signatures”, a <Java parameter declaration list> shall be equivalent to the default Java method signature as determined in Subclause 8.5, “Java routine signature determination”.
- 15) Specifications for Feature J641, “Static fields”:
  - a) Subclause 9.3, “<user-defined type definition>”:
    - i) Without Feature J641, “Static fields”, conforming implementations of this part of ISO/IEC 9075 shall not specify a <static field method spec>.
- 16) Specifications for Feature J651, “SQL/JRT Information Schema”:
  - a) Subclause 13.2, “JARS view”:
    - i) Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA.JARS.
  - b) Subclause 13.3, “METHOD\_SPECIFICATIONS view”:
    - i) Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA . METHOD\_SPECIFICATIONS . EXTERNAL\_NAME or INFORMATION\_SCHEMA . METHOD\_SPECIFICATIONS . IS\_FIELD.
  - c) Subclause 13.6, “USER\_DEFINED\_TYPES view”:
    - i) Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_NAME, INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_LANGUAGE, or INFORMATION\_SCHEMA . UDT\_S . JAVA\_INTERFACE.
  - d) Subclause 13.7, “Short name views”:
    - i) Without Feature J651, “SQL/JRT Information Schema”, conforming SQL language shall not reference INFORMATION\_SCHEMA . METHOD\_SPECS . EXTERNAL\_NAME,

INFORMATION\_SCHEMA . METHOD\_SPECS . IS\_FIELD, INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_NAME, INFORMATION\_SCHEMA . UDT\_S . EXTERNAL\_LANGUAGE, or INFORMATION\_SCHEMA . UDT\_S . JAVA\_INTERFACE.

17) Specifications for Feature J652, “SQL/JRT Usage tables”:

a) Subclause 13.1, “JAR\_JAR\_USAGE view”:

- i) Without Feature J652, “SQL/JRT Usage tables”, conforming SQL language shall not reference INFORMATION\_SCHEMA . JAR\_JAR\_USAGE.

b) Subclause 13.4, “ROUTINE\_JAR\_USAGE view”:

- i) Without Feature J652, “SQL/JRT Usage tables”, conforming SQL language shall not reference INFORMATION\_SCHEMA . ROUTINE\_JAR\_USAGE.

c) Subclause 13.5, “TYPE\_JAR\_USAGE view”:

- i) Without Feature J652, “SQL/JRT Usage tables”, conforming SQL language shall not reference INFORMATION\_SCHEMA . TYPE\_JAR\_USAGE.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002



## Annex B (informative)

### Implementation-defined elements

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

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

- 1) Subclause 4.8.3, “Converting objects between SQL and Java”:
  - a) 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.
- 2) Subclause 4.10, “Privileges”:
  - a) The privileges required to invoke the `SQLJ.INSTALL_JAR`, `SQLJ.REPLACE_JAR`, and `SQLJ.REMOVE_JAR` procedures are implementation-defined.  
NOTE 60 – This is similar to the implementation-defined privileges required for creating a schema.
  - b) 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 <SQLInvoked routine> or the user-name of the current user.
- 3) Subclause 4.11.1, “Deployment descriptor files”:
  - a) An implementation-defined implementor block can be provided in a deployment descriptor file to allow specification of custom install and remove actions.
- 4) Subclause 5.2, “Names and identifiers”:
  - a) The character set supported, and the maximum lengths of the <package identifier>, <class identifier>, <Java field name>, and <Java method name> are implementation-defined.
- 5) Subclause 6.2, “<new specification>”:
  - a) 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.



- 6) Subclause 8.3, “<routine invocation>”:
- a) If validation of the <Java parameter declaration list> has been implementation-defined to be performed by <routine invocation>, then the Syntax Rules of Subclause 8.5, “Java routine signature determination”, are applied with <routine invocation>, a method specification index of 0 (zero), and subject routine *SR*.
  - b) For an external Java routine, let  $CPV_i$  be an implementation-defined non-null value of declared type  $T_i$ .
  - c) The method of execution of a subject Java class’s implementation of `writeObject()` to convert a Java value to an SQL value is implementation-defined.
  - d) The method of execution of a subject Java class’s implementation of `writeSQL()` to convert a Java value to an SQL value is implementation-defined.
  - e) The method of execution of a subject Java class’s implementation of `readObject()` to convert an SQL value to a Java object is implementation-defined.
  - f) The method of execution of a subject Java class’s implementation of `readSQL()` to convert an SQL value to a Java object is implementation-defined.
  - g) If *R* is an external Java routine, then if the JDBC connection object that created any element of *RS* is closed, then the effect is implementation-defined.
  - h) If *R* is an external Java routine, if any element of *RS* is not an object returned by a connection to the current SQL system and SQL session, then the effect is implementation-defined.
  - i) If *R* is an external Java routine, then whether the call of *P* returns update counts as defined in JDBC is implementation-defined.
- 7) Subclause 9.3, “<user-defined type definition>”:
- a) If an <interface using clause> is not explicitly specified, then an implementation-defined <interface specification> is implicit.
  - b) If *UDT* is an external Java data type, then it is implementation-defined whether validation of the explicit or implicit <Java parameter declaration list> is performed by <user-defined type definition> or when the corresponding SQL-invoked method is invoked.
- 8) Subclause 9.4, “<attribute definition>”:
- a) The method of execution of a subject Java class’s implementation of `writeObject()` to convert a Java value to an SQL value is implementation-defined.
  - b) The method of execution of a subject Java class’s implementation of `writeSQL()` to convert a Java value to an SQL value is implementation-defined.
  - c) The method of execution of a subject Java class’s implementation of `readObject()` to convert an SQL value to a Java object is implementation-defined.
  - d) The method of execution of a subject Java class’s implementation of `readSQL()` to convert an SQL value to a Java object is implementation-defined.

- 9) Subclause 9.7, “<SQL-invoked routine>”:
- a) The maximum value of <maximum dynamic result sets> is implementation-defined.
  - b) If *R* is an external Java routine, then it is implementation-defined whether validation of the explicit or implicit <Java parameter declaration list> is performed by <SQL-invoked routine> or when its SQL-invoked routine is invoked.
- 10) Subclause 11.1, “SQLJ.INSTALL\_JAR procedure”:
- a) The maximum length for the CHARACTER VARYING parameters is an implementation-defined integer value.
  - b) The privileges required to invoke the SQLJ.INSTALL\_JAR procedure are implementation-defined.
  - c) The SQLJ.INSTALL\_JAR procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions.
  - d) If an invocation of SQLJ.INSTALL\_JAR raises an exception condition, then the effect on the install actions is implementation-defined.
  - e) The values of the *url* parameter that are valid are implementation-defined, and may include URLs whose format is implementation-defined. If the value of the *url* parameter does not conform to implementation-defined restrictions and does not identify a valid JAR, then an exception condition is raised: *Java DDL — invalid URL*.
- 11) Subclause 11.2, “SQLJ.REPLACE\_JAR procedure”:
- a) The maximum length for the CHARACTER VARYING parameters is an implementation-defined integer value.
  - b) The privileges required to invoke the SQLJ.REPLACE\_JAR procedure are implementation-defined.
  - c) The SQLJ.REPLACE\_JAR procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions.
  - d) The values of the *url* parameter that are valid are implementation-defined, and may include URLs whose format is implementation-defined. If the value of the *url* parameter does not conform to implementation-defined restrictions and does not identify a valid JAR, then an exception condition is raised: *Java DDL — invalid URL*.
- 12) Subclause 11.3, “SQLJ.REMOVE\_JAR procedure”:
- a) The maximum length for the CHARACTER VARYING parameters is an implementation-defined integer value.
  - b) The privileges required to invoke the SQLJ.REMOVE\_JAR procedure are implementation-defined.
  - c) The SQLJ.REMOVE\_JAR procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions.
  - d) If an invocation of SQLJ.REMOVE\_JAR raises an exception condition, then the effect on the remove actions is implementation-defined.

- 13) Subclause 11.4, “SQLJ.ALTER\_JAVA\_PATH procedure”:
- a) The maximum length for the CHARACTER VARYING parameters is an implementation-defined integer value.
  - b) The privileges required to invoke the SQLJ.ALTER\_JAVA\_PATH procedure are implementation-defined.
  - c) The SQLJ.ALTER\_JAVA\_PATH procedure is subject to implementation-defined rules for executing SQL-schema statements within SQL-transactions.
  - d) If an invocation of the SQLJ.ALTER\_JAVA\_PATH procedure raises an exception condition, then effect on the path associated with the JAR is implementation-defined.
- 14) Subclause 12.1.1, “Package java.sql”:
- a) SQL systems that implement this part of ISO/IEC 9075 support the package java.sql, which is the JDBC driver, and all classes required by that package. The other Java packages supplied by SQL systems that implement this part of ISO/IEC 9075 are implementation-defined.
  - b) In an SQL system that implements this part of ISO/IEC 9075, the package java.sql supports the default connection. Other data source URLs that are supported by java.sql are implementation-defined.
- 15) Subclause 12.2, “Deployment descriptor files”:
- a) An <implementor name> is an implementation-defined SQL identifier.
  - b) Whether an <implementor block> with a given <implementor name> contained in an <install actions> (<remove actions>) is interpreted as an install action (remove action) is implementation-defined. That is, an implementation may or may not perform install or remove actions specified by some other implementation.

## Annex C (informative)

### Implementation-dependent elements

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

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

- 1) Subclause 3.3.1, “Specification of built-in procedures”:
  - a) The manner in which built-in procedures are defined is implementation-dependent.
- 2) Subclause 4.8, “User-defined types”:
  - a) The scope and persistence of any modifications to static attributes made during the execution of a Java method is implementation-dependent.
- 3) Subclause 8.3, “<routine invocation>”:
  - a) If  $R$  is an external Java routine, then the scope and persistence of any modifications of class variables made before the completion of any execution of  $P$  is implementation-dependent.
  - b) If the language specifies ADA (respectively C, COBOL, FORTRAN, JAVA, MUMPS, PASCAL, PLI) and  $P$  is not a standard-conforming Ada program (respectively C, COBOL, Fortran, Java, MUMPS, Pascal, PL/I program), then the results of any execution of  $P$  are implementation-dependent.
- 4) Subclause 11.2, “SQLJ.REPLACE\_JAR procedure”:
  - a) The effect of SQLJ.REPLACE\_JAR on currently executing SQL statements that use an SQL routine or structured type whose implementation has been replaced is implementation-dependent.
- 5) Subclause 11.3, “SQLJ.REMOVE\_JAR procedure”:
  - a) The effect of SQLJ.REMOVE\_JAR on currently executing SQL statements that use an SQL routine or structured type whose implementation has been removed is implementation-dependent.

- 6) Subclause 11.4, “SQLJ.ALTER\_JAVA\_PATH procedure”:
  - a) The effect of SQLJ.ALTER\_JAVA\_PATH on SQL statements that have already been prepared or are currently executing is implementation-dependent.

IECNORM.COM : Click to view the full PDF of ISO/IEC 9075-13:2002

## Annex D (informative)

### SQL Feature Taxonomy

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

Table 4, “Feature taxonomy for features outside Core SQL”, contains a taxonomy of the features of the SQL language that are specified in this part of ISO/IEC 9075.

In this table, the first column contains a counter that may be used to quickly locate rows of the table; these values otherwise have no use and are not stable — that is, they are subject to change in future editions of or even Technical Corrigenda to ISO/IEC 9075 without notice.

The “Feature ID” column of Table 4, “Feature taxonomy for features outside Core SQL”, specifies the formal identification of each feature and each subfeature contained in the table. The Feature ID is stable and can be depended on to remain constant. A Feature ID value comprises either a letter and three digits or a letter, three digits, a hyphen, and one or two additional digits. Feature ID values containing a hyphen and additional digits indicate “subfeatures” that help to define complete features, which are in turn indicated by Feature ID values without a hyphen. Only entire features are used to specify the contents of Core SQL and various packages.

**Table 4—Feature taxonomy for features outside Core SQL**

	<b>Feature ID</b>	<b>Feature Name</b>
1	<b>J511</b>	<b>Commands</b> Supports at least one of the following features as SQL statements: Feature J621, “external Java routines” Feature J622, “external Java types” Feature J561, “JAR privileges”
2	<b>J521</b>	<b>JDBC data types</b> Subclause 9.7, “<SQL-invoked routine>”, specifies JDBC data type clauses for <SQL-invoked routine> whose <language clause> specifies JAVA.
3	<b>J531</b>	<b>Deployment</b> Supports at least one of the following features as SQL statements that can be specified in <deployment file>s: Feature J621, “external Java routines” Feature J622, “external Java types” Feature J561, “JAR privileges”
4	<b>J541</b>	<b>SERIALIZABLE</b> Subclause 9.3, “<user-defined type definition>”, specifies SERIALIZABLE

Table 4—Feature taxonomy for features outside Core SQL (Cont.)

	Feature ID	Feature Name
5	<b>J551</b>	<b>SQLDATA</b> Subclause 9.3, “<user-defined type definition>”, specifies SQLDATA
6	<b>J561</b>	<b>JAR privileges</b> Subclause 10.1, “<grant privilege statement>”, specifies GRANT for JARs Subclause 10.3, “<revoke statement>”, specifies REVOKE for JARs Subclause 10.2, “<privileges>”, defines a JAR form of <object name> for granting and revoking USAGE on installed JARs
7	<b>J571</b>	<b>NEW operator</b> Subclause 6.2, “<new specification>”, specifies NEW
8	<b>J581</b>	<b>Output parameters</b> Subclause 9.7, “<SQL-invoked routine>”, specifies output parameter clauses for an external Java routine
9	<b>J591</b>	<b>Overloading</b> Subclause 9.3, “<user-defined type definition>”, specifies one or more <method specification>s whose <method name>s are equivalent to other <method name>s specified by the same <user-defined type definition>
10	<b>J601</b>	<b>SQL-Java paths</b> Subclause 8.2, “<SQL Java path>”, Subclause 11.4, “SQLJ.ALTER_JAVA_PATH procedure”
11	<b>J611</b>	<b>References</b> Subclause 8.3, “<routine invocation>” specifies a <reference expression>; Subclause 9.3, “<user-defined type definition>”, specifies a reference type clause
12	<b>J621</b>	<b>external Java routines</b> Subclause 9.7, “<SQL-invoked routine>”, specifying LANGUAGE JAVA and Subclause 9.9, “<drop routine statement>”, for an external Java routine.
13	<b>J622</b>	<b>external Java types</b> Subclause 9.3, “<user-defined type definition>”, specifying LANGUAGE JAVA, Subclause 9.6, “<drop data type statement>”, for an external Java type, Subclause 9.10, “<user-defined ordering definition>”, for an external Java type, and Subclause 9.11, “<drop user-defined ordering statement>”, for an external Java type.
14	<b>J631</b>	<b>Java signatures</b> Subclause 9.7, “<SQL-invoked routine>”, specifies an external Java routine with <Java parameter declaration list>.
15	<b>J641</b>	<b>Static fields</b> Subclause 9.3, “<user-defined type definition>”, specifies static field methods
16	<b>J651</b>	<b>SQL/JRT Information Schema</b> Subclause 13.2, “JARS view”, Subclause 13.3, “METHOD_SPECIFICATIONS view”, Subclause 13.6, “USER_DEFINED_TYPES view”, Subclause 13.7, “Short name views”,
17	<b>J652</b>	<b>SQL/JRT Usage tables</b> Subclause 13.1, “JAR_JAR_USAGE view”, Subclause 13.4, “ROUTINE_JAR_USAGE view”, Subclause 13.5, “TYPE_JAR_USAGE view”

## Annex E (informative)

### Routines tutorial

#### E.1 Technical components

This part of ISO/IEC 9075 includes the following:

— New built-in procedures.

- `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.

— New built-in schema.

The built-in schema named `SQLJ` is assumed to be in all catalogs of an SQL system that implements the SQL/JRT facility, and to contain all of the built-in procedures of the SQL/JRT facility.

— Extensions of the following SQL statements:

- `CREATE PROCEDURE/FUNCTION` — to specify an SQL name for a Java method.
- `DROP PROCEDURE/FUNCTION` — to delete the SQL name of a Java method.
- `CREATE TYPE` — to specify an SQL name for a Java class.
- `DROP TYPE` — to delete the SQL name of a Java class.
- `GRANT` — to grant the `USAGE` privilege on Java JARs.
- `REVOKE` — to revoke the `USAGE` privilege on Java JARs.

— Conventions for returning values of `OUT` and `INOUT` parameters, and for returning SQL result sets.

— New forms of reference: Qualified references to the fields and methods of columns whose data types are defined on Java classes.



- Additional views and columns in the Information Schema.

## **E.2 Overview**

This tutorial shows a series of example Java classes, indicates how they can be installed, and shows how their static, public methods can be referenced with SQL/JRT facilities in an SQL-environment.

The example Java methods assume an SQL table named EMPS, with the following columns:

- NAME — the employee's name.
- ID — the employee's identification.
- STATE — the state in which the employee is located.
- SALES — the amount of the employee's sales.
- JOBCODE — the job code of the employee.

The table definition is:

```
CREATE TABLE emps (  
  name    VARCHAR(50),  
  id      CHARACTER(5),  
  state   CHARACTER(20),  
  sales   DECIMAL (6,2),  
  jobcode INTEGER );
```

The example classes and methods are:

- `Routines1.region` — A Java method that maps a US state code to a region number. This method doesn't use SQL internally.
- `Routines1.correctStates` — A Java method that performs an SQL UPDATE statement to correct the spelling of *state* codes. The old and new spellings are specified by input-mode parameters.
- `Routines2.bestTwoEmps` — A Java method that determines the top two employees by their sales, and returns the columns of those two employee rows as output-mode parameter values. This method creates an SQL result set and processes it internally.
- `Routines3.orderedEmps` — A Java method that creates an SQL result set consisting of selected employee rows ordered by the sales column, and returns that result set to the client.
- `Over1.isOdd` and `Over2.isOdd` — Contrived Java methods to illustrate overloading rules.
- `Routines4.job1` and `Routines5.job2` — Java methods that return a string value corresponding to an integer jobcode value. These methods illustrate the treatment of null arguments.
- `Routines6.job3` — Another Java method that returns a string value corresponding to an integer jobcode value. This method illustrates the behavior of static Java variables.

Unless otherwise noted, the methods that invoke SQL use JDBC. One of the methods is shown in both a version using JDBC and a version using SQL/OLB. The others could also be coded with SQL/OLB.

It is assumed that the import statements `import java.sql.*;` and `import java.math.*;` have been included in all classes.

### E.3 Example Java methods: `region` and `correctStates`

This clause shows an example Java class, `Routines1`, with two simple methods.

- The `int`-valued static method `region` categorizes 9 states into 3 geographic regions, returning an integer indicating the region associated with a valid state or throwing an exception for invalid states. This method will be called as a function in SQL.
- The `void` method `correctStates` updates the EMPS table to correct spelling errors in the state column. This method will be called as a procedure in SQL.

```
public class Routines1 {
    //An int method that will be called as a function
    public static int region(String s) throws SQLException {
        if (s.equals("MN") || s.equals("VT") || s.equals("NH")) return 1;
        else if (s.equals("FL") || s.equals("GA") || s.equals("AL")) return 2;
        else if (s.equals("CA") || s.equals("AZ") || s.equals("NV")) return 3;
        else throw new SQLException("Invalid state code", "38001");
    }

    //A void method that will be called as a stored procedure
    public static void correctStates (String oldSpelling, String newSpelling)
        throws SQLException {
        Connection conn = DriverManager.getConnection ("jdbc:default:connection");
        PreparedStatement stmt = conn.prepareStatement
            ("UPDATE emps SET state = ? WHERE state = ?");
        stmt.setString(1, newSpelling);
        stmt.setString(2, oldSpelling);
        stmt.executeUpdate();
        stmt.close();
        conn.close();
        return;
    }
}
```

### E.4 Installing `region` and `correctStates` in SQL

The source code for Java classes such as `Routines1` will normally be in one or more Java files (*i.e.*, files with file type “java”). When you compile them (using the `javac` compile command), the resulting code will be in one or more class files (*i.e.*, files with file type “class”). You then typically collect a set of class files into a Java JAR, which is a ZIP-coded collection of files.

To use Java classes in SQL, you load a JAR containing them into the SQL system by calling the SQL `SQLJ.INSTALL_JAR` procedure. The `SQLJ.INSTALL_JAR` procedure is a new built-in SQL procedure that makes the collection of Java classes contained in a specified JAR available for use in the current SQL catalog. For example, assume that you have assembled the above `Routines1` class into a JAR with local file name “`~/classes/Routines1.jar`”:

```
SQLJ.INSTALL_JAR('file:~/classes/Routines1.jar', 'routines1_jar', 0)
```

**E.4 Installing region and correctStates in SQL**

- The first parameter of the `SQLJ.INSTALL_JAR` procedure is a character string specifying the URL of the given JAR. This parameter is never folded to upper case.
- The second parameter of the `SQLJ.INSTALL_JAR` procedure is a character string that will be used as the name of the JAR in the SQL system. The JAR name is an SQL qualified name, and follows SQL conventions for qualified names.

The JAR name that you specify as the second parameter of the `SQLJ.INSTALL_JAR` procedure identifies the JAR within the SQL system. That is, the JAR name that you specify is used only in SQL, and has nothing to do with the contents of the JAR itself. The JAR name is used in the following contexts, which are described in later clauses:

- As a parameter of the `SQLJ.REMOVE_JAR` and `SQLJ.REPLACE_JAR` procedures.
- As a qualifier of Java class names in SQL CREATE PROCEDURE/FUNCTION statements.
- As an operand of the extended SQL GRANT and REVOKE statements.
- As a qualifier of Java class names in SQL CREATE TYPE statements.

The JAR name may also be used in follow-on facilities for downloading JARs from the SQL system.

- JARs can also contain *deployment descriptors*, which specify implicit actions to be taken by the `SQLJ.INSTALL_JAR` and `SQLJ.REMOVE_JAR` procedures. The third parameter of the `SQLJ.INSTALL_JAR` procedure is an integer that specifies whether you do or do not (indicated by non-zero or zero values, respectively) want the `SQLJ.INSTALL_JAR` procedure to execute the actions specified by a deployment descriptor in the JAR. Deployment descriptors are further described in Subclause 12.2, “Deployment descriptor files”.

The name of the `INSTALL_JAR` procedure is qualified with the schema name `SQLJ`. All built-in procedures of the SQL/JRT facility are defined to be contained in that built-in schema. The `SQLJ` schema is assumed to be present in each catalog of an SQL system that implements the SQL/JRT facility.

The first two parameters of `SQLJ.INSTALL_JAR` are character strings, so if you specify them as literals, you will use single quotes, not the double quotes used for SQL delimited identifiers.

The actions of the `SQLJ.INSTALL_JAR` procedure are as follows:

- Obtain the JAR designated by the first parameter.
- Extract the class files that it contains and install them into the current SQL schema.
- Retain a copy of the JAR itself, and associate it with the value of the second parameter.
- If the third parameter is non-zero, then perform the actions specified by the deployment descriptor of the JAR.

After you install a JAR with the `SQLJ.INSTALL_JAR` procedure, you can reference the static methods of the classes contained in that JAR in the CREATE PROCEDURE/FUNCTION statement, as we will describe in the next Subclause.

## E.5 Defining SQL names for region and correctStates

Before you can call a Java method in SQL, you must define an SQL name for it. You do this with new options on the SQL CREATE PROCEDURE/FUNCTION statement. For example:

```
CREATE PROCEDURE correct_states(old CHARACTER(20), new CHARACTER(20))
MODIFIES SQL DATA
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines1_jar:Routines1.correctStates';

CREATE FUNCTION region_of(state CHARACTER(20)) RETURNS INTEGER
NO SQL
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines1_jar:Routines1.region';
```

The CREATE PROCEDURE and CREATE FUNCTION statements specify SQL names and Java method signatures for the Java methods specified in the EXTERNAL NAME clauses. The format of the method names in the external name clause consists of the JAR name that was specified in the SQLJ.INSTALL\_JAR procedure followed by the Java method name, fully qualified with the package name(s) (if any) and class name.

The CREATE PROCEDURE for correct\_states specifies the clause MODIFIES SQL DATA. This indicates that the specified Java method modifies (via INSERT, UPDATE, or DELETE) data in SQL tables. The CREATE FUNCTION for region\_of specifies NO SQL. This indicates that the specified Java method performs no SQL operations.

Other clauses that you can specify are READS SQL DATA, which indicates that the specified Java method reads (through SELECT) data in SQL tables, but does not modify SQL data, and CONTAINS SQL, which indicates that the specified method invokes SQL operations, but neither reads nor modifies SQL data. The alternative CONTAINS SQL is the default.

You use the SQL procedure and function names that you define with such CREATE PROCEDURE/FUNCTION statements as normal SQL procedure and function names:

```
SELECT name, region_of(state) AS region
FROM emps
WHERE region_of(state) = 3;

CALL correct_states ('GEO', 'GA');
```

You can define multiple SQL names for the same Java method:

```
CREATE PROCEDURE state_correction(old CHARACTER(20), new CHARACTER(20))
MODIFIES SQL DATA
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines1_jar:Routines1.correctStates';

CREATE FUNCTION state_region(state CHARACTER(20)) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines1_jar:Routines1.region';
```

The various SQL function and procedure names for a Java method can be used equivalently:

```
SELECT name, state_region(state) AS region
FROM emps
WHERE region_of(state) = 2;

CALL state_correction ('ORE', 'OR');
```

The SQL names are normal 3-part SQL names, and the first two parts of the 3-part names are defaulted as defined in SQL for CREATE PROCEDURE and CREATE FUNCTION statements.

There are other considerations for the CREATE PROCEDURE/FUNCTION statement, dealing with parameter data types, overloaded names, and privileges, which we will discuss in later Subclauses.

## E.6 A Java method with output parameters: bestTwoEmps

The parameters of the `region` and `correctStates` methods are all input-only parameters. This is the normal Java parameter convention.

SQL procedures also support parameters with mode OUT and INOUT. The Java language does not directly have a notion of output parameters. SQL/JRT therefore uses arrays to return output values for parameters of Java methods. That is, if you want an `Integer` parameter to return a value to the caller, you specify the type of that parameter to be `Integer[]`, i.e. an array of `Integer`. Such an array will contain only one element: the input value of the parameter is contained in that element when the method is called, and the method sets the value of that element to the desired output value.

As we will see in the following clauses, this use of arrays for output parameters in the Java methods is visible only to the Java method. When you call such a method as an SQL procedure, you supply normal scalar data items as parameters. The SQL system performs the mapping between those scalar data items and Java arrays implicitly.

The following Java method illustrates the way that you code output parameters in Java. This method, `bestTwoEmps`, returns the name, id, region, and sales of the two employees that have the highest sales in the regions with numbers higher than a parameter value. That is, each of the first 8 parameters is an OUT parameter, and is therefore declared to be an array of the given type.

The following version of the `bestTwoEmps` method uses SQL/OLB for statements that access SQL:

```
public class Routines2 {
    public static void bestTwoEmps (
        String[] n1, String[] id1, int[] r1, BigDecimal[] s1,
        String[] n2, String[] id2, int[] r2, BigDecimal[] s2,
        int regionParm) throws SQLException {
        #sql iterator ByNames (String name, String id, int region, BigDecimal sales);
        n1[0]= "*****"; n2[0]= "*****"; id1[0]= ""; id2[0]= "";
        r1[0]=0; r2[0]=0; s1[0]= new BigDecimal(0); s2[0]= new BigDecimal(0);
        ByNames r;
        try {
            #sql r = {SELECT name, id, region_of(state) AS region, sales
                      FROM emp
                      WHERE region_of(state) > :regionParm
                      AND sales IS NOT NULL
                      ORDER BY sales DESC};

            if (r.next()) {
                n1[0]= r.name();
                id1[0]= r.id();
                r1[0]= r.region();
                s1[0]= r.sales();
            }
            else return;

            if (r.next()) {
                n2[0]= r.name();
                id2[0]= r.id();
                r2[0]= r.region();
                s2[0]= r.sales();
            }
            else return;
        }
    }
}
```

```
    } finally r.close();  
  }  
}
```

Note that since the above Java method uses SQL/OLB for SQL operations, it does not have to explicitly obtain a connection to the SQL system. By default, SQL/OLB executes any SQL contained in a routine in the context of the SQL statement invoking that routine.

For comparison, here's a version of the `bestTwoEmps` method using JDBC instead of SQL/OLB:

```
public class Routines2 {  
    public static void bestTwoEmps (  
        String[ ] n1, String[ ] id1, int[ ] r1, BigDecimal[ ] s1,  
        String[ ] n2, String[ ] id2, int[ ] r2, BigDecimal[ ] s2,  
        int regionParm) throws SQLException {  
        n1[0]= "*****"; n2[0]= "*****"; id1[0]= ""; id2[0]= "";  
        r1[0]=0; r2[0]=0; s1[0]= new BigDecimal(0); s2[0]= new BigDecimal(0);  
        try {  
            Connection conn = DriverManager.getConnection  
                ("jdbc:default:connection");  
            java.sql.PreparedStatement stmt = conn.prepareStatement  
                ("SELECT name, id, region_of(state) AS region, sales  
                FROM emp  
                WHERE region_of(state) > ?  
                AND sales IS NOT NULL  
                ORDER BY sales DESC");  
            stmt.setInt(1, regionParm)  
            ResultSet r = stmt.executeQuery();  
            if (r.next()) {  
                n1[0] = r.getString("name");  
                id1[0] = r.getString("id");  
                r1[0] = r.getInt("region");  
                s1[0] = r.getBigDecimal("sales");  
            }  
            else return;  
            if (r.next()) {  
                n2[0] = r.getString("name");  
                id2[0] = r.getString("id");  
                r2[0] = r.getInt("region");  
                s2[0] = r.getBigDecimal("sales");  
            }  
            else return;  
        } finally { stmt.close(); }  
    }  
}
```

## **E.7 A CREATE PROCEDURE best2 for bestTwoEmps**

Assume that you call the `SQLJ.INSTALL_JAR` procedure for a JAR containing the `Routines2` class with the `bestTwoEmps` method:

```
SQLJ.INSTALL_JAR ('file:~/classes/Routines2.jar', 'routines2_jar', 0)
```



**E.7 A CREATE PROCEDURE best2 for bestTwoEmps**

As indicated previously, in order to call a method such as `bestTwoEmps` in SQL, you must define an SQL name for it, using the CREATE PROCEDURE statement:

```
CREATE PROCEDURE best2 (
    OUT n1 CHARACTER VARYING(50), OUT id1 CHARACTER VARYING(5), OUT r1 INTEGER,
    OUT s1 DECIMAL(6,2),
    OUT n2 CHARACTER VARYING(50), OUT id2 CHARACTER VARYING(5), OUT r2 INTEGER,
    OUT s2 DECIMAL(6,2), region INTEGER)
READS SQL DATA
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines2_jar:Routines2.bestTwoEmps';
```

For parameters that are specified to be OUT or INOUT, the corresponding Java parameter must be an array of the corresponding data type.

**E.8 Calling the best2 procedure**

After you have installed the `Routines2` class in an SQL system and executed the CREATE PROCEDURE for `best2`, you can call the `bestTwoEmps` method as if it were an SQL stored procedure, with normal conventions for OUT parameters. Such a call could be written with embedded SQL, CLI, ODBC, or JDBC. The following is an example of such a call using JDBC:

```
java.sql.CallableStatement stmt = conn.prepareCall(
    "{call best2(?,?,?,?,?,?,?,?)}" );

stmt.registerOutParameter(1, java.sql.Types.STRING);
stmt.registerOutParameter(2, java.sql.Types.STRING);
stmt.registerOutParameter(3, java.sql.Types.INTEGER);
stmt.registerOutParameter(4, java.sql.Types.DECIMAL);
stmt.registerOutParameter(5, java.sql.Types.STRING);
stmt.registerOutParameter(6, java.sql.Types.STRING);
stmt.registerOutParameter(7, java.sql.Types.INTEGER);
stmt.registerOutParameter(8, java.sql.Types.DECIMAL);
stmt.setInt(9, 3);

stmt.executeUpdate();

String n1 = stmt.getString(1);
String id1 = stmt.getString(2);
int r1 = stmt.getInt(3);
BigDecimal s1 = stmt.getBigDecimal(4);
String n2 = stmt.getString(5);
String id2 = stmt.getString(6);
int r2 = stmt.getInt(7);
BigDecimal s2 = stmt.getBigDecimal(8);
```

**E.9 A Java method returning a result set: orderedEmps**

SQL stored procedures can generate SQL result sets as their output. An SQL result set (as defined in JDBC and SQL) is an ordered sequence of SQL rows. SQL result sets aren't processed as normal function result values, but are instead bound to caller-specified iterators or cursors, which are subsequently used to process the rows of the result set.

The following Java method, `orderedEmps`, generates an SQL result set and then returns that result set to the client. Note that the `orderedEmps` method internally generates the result set in the same way as the `bestTwoEmps` method. However, the `bestTwoEmps` method processes the result set within the `bestTwoEmps` method itself, whereas this `orderedEmps` method returns the result set to the client as an SQL result set.

**E.9 A Java method returning a result set: orderedEmps**

To write a Java method that returns a result set to the client, you specify the method to have an additional parameter that is a single-element array of either the Java `ResultSet` class or a class generated by an SQL/OLB iterator declaration (“#sql iterator...”).

The following version of the `orderedEmps` procedure uses SQL/OLB to access the SQL server, and returns the result set as an SQL/OLB iterator, `SalesReport`:

```
// #sql public iterator SalesReport (String name, int region, BigDecimal sales);

public class Routines3 {

    public static void orderedEmps (int regionParm, SalesReport[ ] rs)
        throws SQLException {
        #sql rs[0] = { SELECT name, region_of(state) AS region, sales
                      FROM emp
                      WHERE region_of(state) > :regionParm
                      AND sales IS NOT NULL
                      ORDER BY sales DESC };

        return;
    }
}
```

The `SalesReport` iterator class could be a public static inner class of `Routines3`. However, the above example presumes existence of an “\*.sqlj” file, named `SalesReport.sqlj`, in the same package as `Routines3`, containing the public definition of the `SalesReport` iterator. That is, `SalesReport.sqlj` contains:

```
#sql public iterator SalesReport (String name, int region, BigDecimal sales);
```

Assume, for this example, that both class `Routines3` and the iterator `SalesReport` are defined in a package named `classes`.

For comparison, the following shows `orderedEmps` written using JDBC instead of SQL/OLB.

```
public class Routines3 {

    public static void orderedEmps(int regionParm, ResultSet[ ] rs)
        throws SQLException {

        Connection conn = DriverManager.getConnection ("jdbc:default:connection");
        java.sql.PreparedStatement stmt = conn.prepareStatement
            ("SELECT name, region_of(state) AS region, sales
             FROM emp WHERE region_of(state) > ?
             AND sales IS NOT NULL
             ORDER BY sales DESC");

        stmt.setInt (1, regionParm);
        rs[0] = stmt.executeQuery();
        return;
    }
}
```

The method sets the first element of the `ResultSet[ ]` parameter to reference the Java `ResultSet` containing the SQL result set to be returned. The method does *not* close either the returned `ResultSet` object or the Java statement object that generated the result set. The SQL system will implicitly close both of those objects.

You can call a method such as `orderedEmps` in Java in the normal manner, supplying explicit arguments for both parameters. You can also call it in SQL, as a stored procedure that generates a result set to be processed in the SQL manner. We illustrate how this is done in the following two clauses.



**E.9 A Java method returning a result set: `orderedEmps`**

Each of the above `orderedEmps` examples has a single result set parameter, `rs`, in which you can only return a single result set. You can also specify multiple result set parameters. See Sub-clause 9.7, “<SQL-invoked routine>”.

Note that, in comparison to the prior examples of `bestTwoEmps`, there is no `try...finally` block to close the SQL/OLB iterator or `ResultSet`, `rs[0]`, or the JDBC `PreparedStatement`, `stmt`. For a result set to be returned from a stored procedure it must not be explicitly closed, which means, in the case of JDBC, that the statement executed to generate the result set also must not be explicitly closed.

**E.10 A CREATE PROCEDURE `rankedEmps` for `orderedEmps`**

Assume that you call the `SQLJ.INSTALL_JAR` procedure for a JAR containing the `Routines3` class with the `orderedEmps` method:

```
SQLJ.INSTALL_JAR( 'file:~/classes/Routines3.jar', 'routines3_jar', 0)
```

As with previous methods, you will now need to define an SQL name for the `orderedEmps` method before you can call it as an SQL procedure. As above, you will do this with a `CREATE PROCEDURE` statement that specifies an `EXTERNAL...LANGUAGE JAVA` clause to reference the `orderedEmps` method. The following is an example `CREATE PROCEDURE...DYNAMIC RESULT SETS` for the above `orderedEmps` method:

```
CREATE PROCEDURE ranked_emp (region INTEGER)
  READS SQL DATA
  DYNAMIC RESULT SETS 1
  LANGUAGE JAVA PARAMETER STYLE JAVA
  EXTERNAL NAME 'routines3_jar:classes.Routines3.orderedEmps';
```

A `CREATE PROCEDURE` statement for a Java method that generates SQL result sets has the following characteristics:

- The `DYNAMIC RESULT SETS` clause indicates that the procedure generates one or more result sets. The integer specified in the `DYNAMIC RESULT SETS` clause is the maximum number of result sets that the procedure will generate. If an execution generates more than this number of result sets, a warning will be issued, and only the specified number of result sets will be returned.
- The SQL signature specifies only the parameters that the caller explicitly supplies.
- The specified Java method actually has one or more additional, trailing parameters, whose data types must be a Java array of either `java.sql.ResultSet` or an implementation of `sqlj.runtime.ResultSetIterator`.

The above `CREATE PROCEDURE` statement could be used to reference either an SQL/OLB-based or JDBC-based version of `Routines3.orderedEmps`. When it is necessary to choose a particular implementation, the Java method signature of the desired Java method must be explicitly stated. For the SQL/OLB-based `orderedEmps`:

```
CREATE PROCEDURE ranked_emp (region INTEGER)
  READS SQL DATA
  DYNAMIC RESULT SETS 1
  LANGUAGE JAVA PARAMETER STYLE JAVA
  EXTERNAL NAME
    'routines3_jar:classes.Routines3.orderedEmps(int, classes.SalesReport[]);'
```

**E.10 A CREATE PROCEDURE rankedEmps for orderedEmps**

And, for the JDBC-based orderedEmps:

```
CREATE PROCEDURE ranked_emp (region INTEGER)
  READS SQL DATA
  DYNAMIC RESULT SETS 1
  LANGUAGE JAVA PARAMETER STYLE JAVA
  EXTERNAL NAME
    'routines3_jar:classes.Routines3.orderedEmps(int, java.sql.ResultSet[])';
```

The only difference in the above CREATE PROCEDURE ranked\_emp statements is in the Java method signature's description of the dynamic result set returned. In both cases, a fully qualified class name is provided for, respectively, the SQL/OLB iterator (remember that SalesReport is in the package named classes) and the JDBC result set.

The next clause will show an example invocation of this procedure.

**E.11 Calling the rankedEmps procedure**

After you have installed the Routines3 class in an SQL system and executed the CREATE PROCEDURE for rankedEmps, you can call the rankedEmps procedure as if it were an SQL stored procedure. Such a call could be written with any facility that defines mechanisms for processing SQL result sets — that is, SQL/CLI, JDBC, and SQL/OLB. The following is an example of such a call using JDBC:

```
java.sql.CallableStatement stmt = conn.prepareCall( "{call ranked_emp(?)}");
stmt.setInt(1, 3);
ResultSet rs = stmt.executeQuery();
while (rs.next()) {
  String name = rs.getString(1);
  int region = rs.getInt(2);
  BigDecimal sales = rs.getBigDecimal(3);
  System.out.print("Name = " + name);
  System.out.print("Region = " + region);
  System.out.print("Sales = " + sales);
  System.out.println();
}
```

Note that the call of the ranked\_emp procedure supplies only the single parameter that was declared in the CREATE PROCEDURE statement. The SQL system then implicitly supplies, as applicable, a parameter that is an empty array of ResultSet or an empty array of classes.SalesReport, and calls the Java method. That Java method assigns the output result set or iterator to the array parameter. And, when the Java method completes, the SQL system returns the result set or iterator in that output array element as an SQL result set.

**E.12 Overloading Java method names and SQL names**

When you use CREATE PROCEDURE/FUNCTION statements to specify SQL names for Java methods, the SQL names can be overloaded. That is, you can specify the same SQL name in multiple CREATE PROCEDURE/FUNCTION statements. Note that support for such SQL overloading is an optional feature.

Consider the following Java classes and methods. These are contrived routines intended only to illustrate overloading, and we won't show the routine bodies.

**E.12 Overloading Java method names and SQL names**

```

public class Over1 {
    public static int isOdd (int i) {...};
    public static int isOdd (float f) {...};
    public static int testOdd (double d) {...};
}

public class Over2 {
    public static int isOdd (java.sql.Timestamp t) {...};
    public static int oddDateTime (java.sql.Date d) {...};
    public static int oddDateTime (java.sql.Time t) {...};
}

```

Note that the `isOdd` method name is overloaded in the `Over1` class, and the `oddDateTime` method name is overloaded in the `Over2` class.

Assume that the above classes are in a JAR `~/classes/Over.jar`, which you install:

```
SQLJ.INSTALL_JAR ('file:~/classes/Over.jar', 'over_jar', 0)
```

To reference these methods in SQL, you will of course need to specify SQL names for them with `CREATE FUNCTION` statements. These `CREATE FUNCTION` statements can specify SQL names that are overloaded. The overloading of the SQL names is completely separate from the overloading in the Java names. This is illustrated in the following.

Recall that you can specify the same Java method in multiple `CREATE PROCEDURE/FUNCTION` statements.

```

CREATE FUNCTION odd (INTEGER) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over1.isOdd';
CREATE FUNCTION odd (REAL) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over1.isOdd';
CREATE FUNCTION odd (DOUBLE PRECISION) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over1.testOdd';
CREATE FUNCTION odd (TIMESTAMP) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over2.isOdd';
CREATE FUNCTION odd (DATE) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over2.oddDateTime';
CREATE FUNCTION odd (TIME) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over2.oddDateTime';
CREATE FUNCTION is_odd (INTEGER) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over1.isOdd';
CREATE FUNCTION test_odd (REAL) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Over1.isOdd';

```

Note the following characteristics of these `CREATE FUNCTION` statements:

- The SQL name `odd` is defined on the two `isOdd` methods and the `testOdd` method of `Over1`, and also the `isOdd` method and two `oddDateTime` methods of `Over2`. That is, the SQL name `odd` *spans* both overloaded and non-overloaded Java names.
- The SQL names `is_odd` and `test_odd` are defined on the two `isOdd` methods of `Over1`. That is, those two different SQL names are defined on the same Java name.

**E.12 Overloading Java method names and SQL names**

The rules governing overloading are those of the SQL language as defined in Subclause 11.49, "SQL-invoked routine", in ISO/IEC 9075-2, and in Subclause 10.4, "<routine invocation>", in ISO/IEC 9075-2. This includes:

- Rules governing what parameter combinations can be overloaded. That is, the legality (or not) of the following CREATE statements is determined by SQL language rules:

```
CREATE FUNCTION is_odd (INTEGER) RETURNS INTEGER...
```

```
CREATE FUNCTION is_odd (SMALLINT) RETURNS INTEGER...
```

```
CREATE PROCEDURE is_odd (SMALLINT) ...
```

- Rules governing the resolution of calls using overloaded SQL names. That is, the determination of which Java method is called by "odd(x)" for some data item "x" is determined by SQL language rules.

The EXTERNAL NAME clauses of the above CREATE FUNCTION statements specify only the JAR name and method name of the Java method. For example:

```
CREATE FUNCTION odd (INTEGER) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Overl.isOdd';
```

You can also include the Java method signature (*i.e.*, a list of the parameter data types) of a method in the EXTERNAL NAME clause. For example:

```
CREATE FUNCTION odd (INTEGER) RETURNS INTEGER
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'over_jar:Overl.isOdd (int)';
```

The group of eight example CREATE FUNCTION statements, shown earlier in this clause, do not require Java method signatures, but you can include them for clarity. Subclause E.14, "Java method signatures in the CREATE statements", describes cases where the Java method signature is required.

**E.13 Java main methods**

The Java Language Specification places special requirements on any method named `main`. A method named `main` is required to have the following Java method signature:

```
public static void main (String[ ]);
```

If you specify a Java method named `main` in an SQL CREATE PROCEDURE...EXTERNAL statement, then that Java method must have the above Java method signature. The signature of the SQL procedure can either be:

- A single parameter that is an array of CHARACTER or CHARACTER VARYING. That array is passed to the Java method as the String array parameter. Note: This SQL method signature can only be used in SQL systems that support array data types in SQL.
- Zero or more parameters, each of which is CHARACTER or CHARACTER VARYING. Those *N* parameters are passed to the Java method as a single *N* element array of String.

## E.14 Java method signatures in the CREATE statements

Consider the following method, `job1`, which has an integer parameter and returns a `String` with the job corresponding with a `jobcode` value:

```
public class Routines4 {
    //A String method that will be called as a function
    public static String job1 (Integer jc) throws SQLException {
        if (jc == 1) return "Admin";
        else if (jc == 2) return "Sales";
        else if (jc == 3) return "Clerk";
    else if (jc == null) return null;
        else return "unknown jobcode";
    }
}
```

Note that we suffix the method name with a “1” in anticipation of subsequent variants of the method.

Assume that you install this class in SQL:

```
SQLJ.INSTALL_JAR ('file:~/classes/Routines4.jar', 'routines4_jar', 0)
```

You might want to specify an SQL function `job_of1` defined on the `job1` method:

```
CREATE FUNCTION job_of1(jc INTEGER) RETURNS CHARACTER VARYING(20)
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines4_jar:Routines4.job1';
```

However, as written above, this CREATE statement is not valid. Note that the data type of the parameter of the Java method `job1` is `Java Integer` (which is short for `java.lang.Integer`), and we have specified the SQL data type `INTEGER` for the corresponding parameter of the SQL `job_of1` function. However, the detailed rules (see Subclause 9.7, “<SQL-invoked routine>” for the external Java form of the SQL CREATE PROCEDURE/FUNCTION statement specifies that the default Java parameter data type for an SQL `INTEGER` parameter is the `Java int` data type, not the `Java Integer` data type. Subclause E.15, “Null argument values and the RETURNS NULL clause”, describes some reasons why you may want to specify `Java Integer` rather than `Java int`.

If you want to specify an SQL CREATE PROCEDURE/FUNCTION statement for a Java method whose parameter data types include Java types differing from their default Java types, then you specify those data types in a Java method signature in the CREATE statement. This Java method signature is written after the Java method name in the EXTERNAL NAME clause. For example, the above CREATE statement for the `job1` method would be written as:

```
CREATE FUNCTION job_of1(jc INTEGER) RETURNS CHARACTER VARYING(20)
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines4_jar:Routines4.job1(java.lang.Integer)';
```

If you specify data types in the Java method signature of a CREATE statement that specifies DYNAMIC RESULT SETS, then you must include the implicit trailing result set or iterator parameters in that Java method signature. You do not, however, include those trailing parameters in the SQL signature. For example, you would write the CREATE of Subclause E.10, “A CREATE PROCEDURE `rankedEmps` for `orderedEmps`”, as follows:

```
CREATE PROCEDURE ranked_emp (region INTEGER)
READS SQL DATA
DYNAMIC RESULT SETS 1
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines3_jar:Routines3.orderedEmps (int, java.sql.ResultSet[ ]);
```

See Subclause 9.7, “<SQL-invoked routine>”.

## E.15 Null argument values and the RETURNS NULL clause

Consider the Java method `job1` and the corresponding SQL function `job_of1`, which we defined in Subclause E.14, “Java method signatures in the CREATE statements”.

You can call the SQL function `job_of1` in SQL statements such as the following:

```
SELECT name, job_of1(jobcode)
FROM emps
WHERE job_of1(jobcode) <> 'Admin';
```

Suppose that a row of the EMPS table has a null value in the JOBCODE column. Note that the Java data type of the parameter of the `job1` method is Java Integer (that is, `java.lang.Integer`). The Java Integer data type is a class, rather than a scalar data type, so its values include both numeric values, and also the null reference value. When an SQL null value is passed as an argument to a Java parameter whose data type is a Java class, the null SQL value is passed as a Java null reference. Such a null reference can be tested within the Java method, as shown in `Routines4.job1`.

Now consider the following similar method, which specifies its parameter data type to be the Java scalar data type `int`, rather than the Java class `Integer`.

```
public class Routines5 {
    //A String method that will be called as a function
    public static String job2 (int jc)
        throws SQLException {
        if (jc == 1) return "Admin";
        else if (jc == 2) return "Sales";
        else if (jc == 3) return "Clerk";
        else return "unknown jobcode";
    }
}
```

Assume that you install this class in SQL:

```
SQLJ.INSTALL_JAR( 'file:~/classes/Routines5.jar', 'routines5_jar', 0)
CREATE FUNCTION job_of2 (jc INTEGER) RETURNS CHARACTER VARYING(20)
LANGUAGE JAVA PARAMETER STYLE JAVA
EXTERNAL NAME 'routines5_jar:Routines5.job2';
```

You can then call the SQL function `job_of2` in SQL statements such as the following:

```
SELECT name, job_of2(jobcode)
FROM emps
WHERE job_of2(jobcode) <> 'Admin';
```

When this SELECT statement encounters a row of the EMPS table in which the JOBCODE column is null, the effect of the null value on the call(s) of the `job_of2` function is different than for the previous `job_of` function. The `job_of2` function is defined on the method `Routines5.job2`, whose parameter has the scalar data type `int`, rather than the class data type `java.lang.Integer`. The Java `int` data type (and other Java scalar data types) has no null reference value, and no other representation of a null value. Therefore, if the `job2` method is invoked with a null SQL value, then an exception condition is raised.

To summarize:

- The following Java data types have null reference values, and can accommodate SQL arguments that are null:



**E.15 Null argument values and the RETURNS NULL clause**

```
java.lang.String, java.math.BigDecimal, byte[] , java.sql.Date, java.sql.Time,
java.sql.Timestamp, java.lang.Double, java.lang.Float, java.lang.Integer, java.lang.Short,
java.lang.Long, java.lang.Boolean
```

- The following Java data types are scalar data types that cannot accommodate nulls. An exception condition will be raised if an argument value passed as such a parameter data type is null:

```
boolean, byte, short, int, long, float, double
```

The exception condition that is raised when you attempt to pass a null argument to a Java parameter that is a non-nullable data type is analogous to the traditional SQL exception condition that is raised when you attempt to FETCH or SELECT a null column value into a host variable for which you did not specify a null indicator variable. In both cases, the “receiving” parameter or variable is unable to accommodate the actual run-time null value, so an exception condition is raised.

When you code Java methods specifically for use in SQL, you will probably tend to specify Java parameter data types that are the nullable Java data types. You may, however, also want to use Java methods in SQL that were not coded for use in SQL, and that are more likely to specify Java parameter data types that are the scalar (non-nullable) Java data types.

You can call such functions in contexts where null values will occur by invoking them conditionally, *e.g.*, in CASE expressions. For example:

```
SELECT name,
       CASE
         WHEN jobcode IS NOT NULL THEN job_of2 (jobcode)
         ELSE NULL
       END
FROM emps
WHERE CASE
       WHEN jobcode IS NOT NULL THEN job_of2 (jobcode)
       ELSE NULL
     END <> 'Administrator';
```

You can also make such CASE expressions implicit, by specifying the RETURNS NULL ON NULL INPUT option in the CREATE FUNCTION statement:

```
CREATE FUNCTION job_of22 (jc INTEGER) RETURNS CHARACTER VARYING(20)
  RETURNS NULL ON NULL INPUT
  LANGUAGE JAVA PARAMETER STYLE JAVA
  EXTERNAL NAME 'routines5_jar:Routines5.job2';
```

When an SQL function is called whose CREATE FUNCTION statement specifies RETURNS NULL ON NULL INPUT, then if the runtime value of any argument is null, the result of the function call is set to null, and the function itself is not invoked.

The following SELECT statement invokes the job\_of22 function.

```
SELECT name, job_of22(jobcode)
FROM emps
WHERE job_of22(jobcode) <> 'Administrator';
```

This SELECT is equivalent to the previous SELECT that invokes the job\_of2 function within CASE expressions. That is, the RETURNS NULL ON NULL INPUT clause in the CREATE FUNCTION statement for job\_of22 makes the null-testing CASE expressions implicit.

The RETURNS NULL ON NULL INPUT option applies to *all* of the parameters of the function, not just to the parameters whose Java data type is not nullable.

**E.15 Null argument values and the RETURNS NULL clause**

The convention that the RETURNS NULL ON NULL INPUT option defines for a function is the same convention that is followed for most built-in SQL functions and operators: if any operand is null, then the value of the operation is null.

The alternative to the RETURNS NULL ON NULL INPUT clause is CALLED ON NULL INPUT, which is the default.

You can specify the same Java method in multiple CREATE FUNCTION statements (*i.e.*, defining SQL synonyms), and those CREATE FUNCTION statements can either specify RETURNS NULL ON NULL INPUT or CALLED ON NULL INPUT, as illustrated by the above job\_of2 and job\_of22.

If you create multiple SQL functions named job\_of22 (with different numbers and/or types of parameters), you can specify (or default to) CALLED ON NULL INPUT in some of the CREATE FUNCTION job\_of22 statements, and specify RETURNS NULL ON NULL INPUT in others. The actions of the RETURNS NULL ON NULL INPUT clause are taken after overloading resolution has been done and a particular CREATE FUNCTION statement has been identified.

The RETURNS NULL ON NULL INPUT and CALLED ON NULL INPUT clauses can only be specified in CREATE FUNCTION statements, that is, not in CREATE PROCEDURE statements. This is because there is no equivalent conditional treatment of procedure calls that would be as generally useful.

**E.16 Static variables**

Java static methods can be contained in Java classes that have static variables, and, in Java, static methods can both reference and set static variables. For example:

```
public class Routines6 {
    static String jobs;
    public static void setJobs (String js) throws SQLException {jobs=js;}
    public static String job3(int jc) throws SQLException {
        if (jc < 1 || jc * 5 > length(jobs,41) return "Invalid jobcode";
        else return jobs.substring(5*(jc-1), 5*jc);
    }
}
```

Assume that you install this class in an SQL system:

```
SQLJ.INSTALL_JAR('file:~/classes/Routines6.jar', 'routines6_jar', 0);
```

The class Routines6 has a static variable jobs, which is set by the static method setJobs and referenced by the static method job3. A class such as Routines6 that dynamically modifies the values of static variables is well-defined in Java, and can be useful. However, when such a class is installed in an SQL system, and the methods setJobs and job3 are defined as SQL procedures and functions (<SQL-invoked routine>), the scope of the assignments to the static variable jobs is implementation-dependent. That is, the scope of that variable is not specified, and is likely to differ across implementations (and possibly across the releases of a given implementation).

For example:

```
CREATE PROCEDURE set_jobs (js CHARACTER VARYING(100))
    LANGUAGE JAVA PARAMETER STYLE JAVA
    EXTERNAL NAME 'routines6_jar:Routines6.setJobs';

CREATE FUNCTION job_of3 (jc integer) RETURNS CHARACTER VARYING(20)
    RETURNS NULL ON NULL INPUT
    LANGUAGE JAVA PARAMETER STYLE JAVA
    EXTERNAL NAME 'routines6_jar:Routines6.job3';
```



**E.16 Static variables**

```
CALL set_jobs ('AdminSalesClerk');

SELECT name, job_of3 (jobcode)
FROM emps
WHERE job_of3(jobcode) <> 'Admin';
```

This appears to be a straightforward use of the `Routines6` class in SQL. The call of `set_jobs` specifies a list of job code values, which a user might reasonably assume is “cached” by the SQL-environment and used in subsequent calls of `job_of3`. However, since the scope of the static variable `jobs` in the SQL environment is implementation-dependent, the answers to the following questions regarding the values passed to the `set_jobs` procedure are likely to differ across implementations:

- Is the `set_jobs` value visible only to the current session? Or also to concurrent sessions and to later non-concurrent sessions?
- Does the `set_jobs` value persist across a COMMIT? Is it reset by a ROLLBACK?

The implication of this uncertainty is that you should not use classes that assign to static variables in SQL. Note, however, that such assignments will not (necessarily) be detected by the SQL implementation, either when you CREATE PROCEDURE/FUNCTION or when you call a routine.

You can prevent assignments to static variables in Java by declaring them with the `final` property.

**E.17 Dropping SQL names of Java methods**

After you have created SQL procedure or function names for Java methods, you can drop those SQL names with a normal SQL DROP statement:

```
DROP FUNCTION region RESTRICT;
```

A DROP statement has no effect on the Java method (or class) on which the SQL name was defined. Dropping an SQL procedure or function implicitly revokes any granted privileges for that routine.

**E.18 Removing Java classes from SQL**

You can completely uninstall a JAR with the `SQLJ.REMOVE_JAR` procedure. For example:

```
SQLJ.REMOVE_JAR ('routines_jar', 0);
```

As noted earlier, JARS can contain *deployment descriptors*, which specify implicit actions to be taken by the `SQLJ.INSTALL_JAR` and `SQLJ.REMOVE_JAR` procedures. The second parameter is an integer that specifies whether you do or do not (indicated by non-zero or zero values, respectively) want the `SQLJ.REMOVE_JAR` procedure to execute the actions specified by a deployment descriptor in the JAR. Deployment descriptors are further described in Subclause 12.2, “Deployment descriptor files”.

After the `SQLJ.REMOVE_JAR` procedure performs any actions specified by the JAR’s deployment descriptor file(s), there must be no remaining SQL procedure or function whose external name references any method of any class in the specified JAR. Any such remaining SQL procedures or functions must be explicitly dropped before the `SQLJ.REMOVE_JAR` procedure will be able to complete successfully.

## E.19 Replacing Java classes in SQL

Assume that you have installed a Java JAR in SQL, and you want to replace some or all of the contained classes, *e.g.*, to correct or improve them. You can do this by using the `SQLJ.REMOVE_JAR` procedure to remove the current JAR, and then using the `SQLJ.INSTALL_JAR` procedure to install the new version. However, you will probably have executed one or more SQL DDL statements that depend on the methods of the classes that you want to replace. That is, you may have executed one or more of the following DDL operations:

- `CREATE PROCEDURE/FUNCTION` statements referencing the classes.
- `GRANT` statements referencing those SQL procedures and functions.
- `CREATE PROCEDURE/FUNCTION` statements for SQL procedures and functions that invoke those SQL procedures and functions.
- `CREATE VIEW/TABLE` statements for SQL views and tables that invoke those SQL procedures and functions.

The rules for the `SQLJ.REMOVE_JAR` procedure require that you drop all SQL procedure/functions that directly reference methods of a class before you can remove the JAR containing the class. And, SQL rules for `RESTRICT`, as specified in the SQL <drop routine statement>, require that you drop all SQL objects (tables, views, SQL-server modules, and routines whose bodies are written in SQL) that invoke a procedure/function before you drop the procedure/function.

Thus, if you use the `SQLJ.REMOVE_JAR` and `SQLJ.INSTALL_JAR` procedures to replace a JAR, you will have to drop the SQL objects that directly or indirectly depend on the methods of the classes in the JAR, and then re-create those items.

The `SQLJ.REPLACE_JAR` procedure avoids this requirement, by performing an instantaneous *remove* and *install*, with suitable validity checks. You can therefore call the `SQLJ.REPLACE_JAR` procedure without first dropping the dependent SQL objects.

For example, in Subclause E.4, “Installing region and correctStates in SQL”, we installed the class of `Routines1` with the following statement:

```
SQLJ.INSTALL_JAR( 'file:~/classes/Routines1.jar', 'routines1_jar', 0)
```

You can replace that JAR with a statement such as:

```
SQLJ.REPLACE_JAR( 'file:~/revised_classes/Routines1.jar', 'routines1_jar')
```

Note that the JAR name must be the same. It identifies the existing JAR, and will subsequently identify the replacement JAR. The URL of the replacement JAR can be the same as or different from the URL of the original JAR.

In the general case, there will be classes in the old JAR that are not in the new JAR, classes that are in both JARs, and classes that are in the new JAR and not in the old JAR. These are referred to respectively as unmatched old classes, matching old/new classes, and unmatched new classes.

The validity requirements on the replacement JAR are:

- There must be no SQL procedure or function whose routine descriptor's <external routine name> specified an <external Java reference string> that references any method of any unmatched old class (since all unmatched old classes will be removed).
- Any `CREATE PROCEDURE/FUNCTION` statement that references a method of a matching class must be a valid statement for the new class.