

Oracle8*i*

Java Developer's Guide

Release 3 (8.1.7)

July 2000

Part No. A83728-01

ORACLE®

Java Developer's Guide, Release 3 (8.1.7)

Part No. A83728-01

Copyright © 1996, 2000, Oracle Corporation. All rights reserved.

Primary Author: Sheryl Maring

Contributors: Steve Harris, Ellen Barnes, Peter Benson, Greg Colvin, Bill Courington, Matthieu Devin, Jim Haungs, Hal Hildebrand, Mark Jungerman, Susan Kraft, Thomas Kurian, Scott Meyer, Tom Portfolio, Dave Rosenberg, Jerry Schwarz, Harlan Sexton, Tim Smith, David Unietis, Brian Wright.

The Programs (which include both the software and documentation) contain proprietary information of Oracle Corporation; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. Oracle Corporation does not warrant that this document is error free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of Oracle Corporation.

If the Programs are delivered to the U.S. Government or anyone licensing or using the programs on behalf of the U.S. Government, the following notice is applicable:

Restricted Rights Notice Programs delivered subject to the DOD FAR Supplement are "commercial computer software" and use, duplication, and disclosure of the Programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, Programs delivered subject to the Federal Acquisition Regulations are "restricted computer software" and use, duplication, and disclosure of the Programs shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software - Restricted Rights (June, 1987). Oracle Corporation, 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and Oracle Corporation disclaims liability for any damages caused by such use of the Programs.

Oracle is a registered trademark, and Oracle products mentioned herein are trademarks or registered trademarks of Oracle Corporation. All other company or product names mentioned are used for identification purposes only and may be trademarks of their respective owners.

Contents

Send Us Your Comments	ix
Preface.....	xi
1 Introduction to Java in Oracle8i	
Contents.....	1-2
Overview of Java.....	1-2
Java and Object-Oriented Programming Terminology.....	1-2
Classes.....	1-2
Attributes.....	1-4
Methods.....	1-5
Class Hierarchy.....	1-5
Interfaces.....	1-6
Polymorphism.....	1-7
The Java Virtual Machine (JVM).....	1-8
Key Features of the Java Language.....	1-11
Why Use Java in Oracle8i?	1-12
Multithreading.....	1-13
Automated Storage Management	1-14
Footprint	1-14
Performance.....	1-15
How Native Compilers Improve Performance.....	1-15
Dynamic Class Loading.....	1-16
Oracle's Java Application Strategy	1-17
Java Stored Procedures	1-18

PL/SQL Integration and Oracle RDBMS Functionality	1-18
JDBC Drivers	1-19
SQLJ – Embedded SQL in Java.....	1-19
Dynamic HTML Invoking Java	1-20
Servlets	1-20
JavaServerPages.....	1-20
Distributed Application Development.....	1-21
Using EJB Components	1-21
Development Tools	1-22
Overview of Oracle8i Java Documentation	1-22

2 Writing Java Applications on Oracle8i

Overview	2-2
Terminology	2-2
Database Sessions Imposed on Java Applications.....	2-3
Session Lifetime	2-5
Java Supported APIs	2-5
Execution Control	2-6
Migrating from JDK 1.1 to Java 2.....	2-7
Your Development Environment.....	2-7
JDBC 2.0	2-8
Server Applications Using JDBC 2.0.....	2-9
Clients Using JDBC 2.0	2-10
Java 2 Security	2-10
Java 2 ORB APIs.....	2-10
JNDI Lookup.....	2-11
Aurora ORB Interface	2-11
CORBA ORB Interface	2-13
Backwards Compatibility for 8.1.5 CORBA and EJB Applications.....	2-14
Java Code, Binaries, and Resources Storage.....	2-14
Preparing Java Class Methods for Execution	2-15
Compiling Java Classes.....	2-15
Compiling Source through javac.....	2-15
Compiling Source through loadjava.....	2-15
Compiling Source at Runtime	2-16

Specifying Compiler Options	2-16
Automatic Recompilation	2-19
Resolving Class Dependencies	2-19
Allowing References to Non-Existent Classes	2-21
ByteCode Verifier	2-22
Loading Classes	2-23
Two Definitions of the Same Class	2-25
Need Database Privileges and JVM Permissions	2-26
Loading JAR or ZIP Files.....	2-26
How to Grant Execute Rights	2-27
Checking Java Uploads.....	2-28
Object Name and Type	2-29
Status	2-30
Publishing.....	2-30
User Interfaces on the Server	2-31
Shortened Class Names	2-32
Class.forName() on JServer	2-33
Supply the ClassLoader in Class.forName	2-34
Supply Class and Schema Names to classForNameAndSchema	2-35
Supply Class and Schema Names to lookupClass.....	2-36
Supply Class and Schema Names when Serializing	2-36
Class.forName Example	2-37
Managing Your Operating System Resources	2-38
Overview of Operating System Resources	2-38
Operating System Resource Access	2-39
Operating System Resource Lifetime	2-39
Garbage Collection and Operating System Resources.....	2-40
Operating System Resources Affected Across Calls	2-41
Sockets.....	2-43
Threading in JServer	2-44
Thread Lifecycle	2-45

3 Invoking Java in the Database

Overview	3-2
Invoking Java Methods	3-3

Utilizing Java Stored Procedures	3-3
Utilizing Distributed Objects With CORBA and EJB	3-6
IIOP Transport	3-7
Naming	3-7
Creating and Deploying Enterprise JavaBeans	3-8
Using an EJB.....	3-9
Session Shell	3-9
Utilizing Remote Method Invocation (RMI).....	3-10
Utilizing Java Native Interface (JNI) Support	3-10
Utilizing SQLJ and JDBC for Querying Database.....	3-11
JDBC	3-11
SQLJ.....	3-11
An Example Comparing JDBC and SQLJ	3-12
Complete SQLJ Example	3-13
SQLJ Strong Typing Paradigm.....	3-15
Translating a SQLJ Program	3-16
Running a SQLJ Program in the Server.....	3-16
Converting a Client Application to Run in the Server.....	3-17
Interacting with PL/SQL.....	3-17
Debugging Server Applications	3-18
1. Prepare the Code for Debugging	3-19
2. Start the Debug Proxy.....	3-20
3. Starting, Stopping, and Restarting the Debug Agent.....	3-21
OracleAgent Class	3-22
4. Connecting a Debugger	3-22
How To Tell You Are Executing in the Server	3-24
Redirecting Output on the Server	3-24

4 Java Installation and Configuration

Initializing a Java-Enabled Database.....	4-2
Manual Install	4-2
Requirements	4-3
Package DBMS_JAVA.....	4-3
Configuring JServer.....	4-6
Java Stored Procedure Configuration.....	4-7

Enterprise JavaBeans and CORBA Configuration.....	4-8
Enabling the Java Client	4-8
1. Install JDK on the Client.....	4-8
2. Set up CLASSPATH.....	4-9
Basic Included JAR files:.....	4-9
JAR Files Necessary for JDK 1.1 Clients.....	4-10
JAR Files Necessary for Java 2 Clients.....	4-10
JAR Files Included for Clients that use SQLJ.....	4-10
JAR Files Included for Clients that use JSP.....	4-11
3. Verify the Port/SID.....	4-11
4. Test Install with Samples.....	4-11

5 Security For Oracle8i Java Applications

Network Connection Security.....	5-2
Database Contents and JVM Security.....	5-2
Java 2 Security.....	5-3
Setting Permissions.....	5-6
Fine-Grain Definition for Each Permission.....	5-7
Acquiring Administrative Permission to Update Policy Table.....	5-12
Creating Permissions.....	5-14
Enabling or Disabling Permissions.....	5-17
Permission Types.....	5-19
Initial Permission Grants.....	5-21
General Permission Definition Assigned to Roles.....	5-24
Debugging Permissions.....	5-25
Permission for Loading Classes.....	5-25

6 Oracle8i Java Application Performance

Natively Compiled Code.....	6-2
JServer Accelerator Overview.....	6-3
JServer Core Java Class Libraries.....	6-5
Natively Compiling Java Application Class Libraries.....	6-5
Installation Requirements.....	6-5
Executing JServer Accelerator.....	6-6
Java Memory Usage.....	6-7

Configuring Memory Initialization Parameters.....	6-8
Java Pool Memory.....	6-9
Displaying Used Amounts of Java Pool Memory.....	6-10
Correcting Out of Memory Errors.....	6-11
Running out of memory while compiling	6-11
Running out of memory while loading.....	6-11
End-of-Call Migration	6-12
Oracle-Specific Support for End-of-Call Optimization	6-13
Memory Profiling Utility	6-17
How MemStat Works.....	6-18
Using MemStat.....	6-19
MemStat Permissions	6-20
The MemStat Report Format.....	6-21
Sample Output	6-22

Glossary

Index

Send Us Your Comments

Oracle8i Java Developer's Guide, Release 3 (8.1.7)

Part No. A83728-01

Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the chapter, section, and page number (if available). You can send comments to us in the following ways:

- Electronic mail — jpgcomnt@us.oracle.com
- FAX - 650-506-7225. Attn: Java Platform Group, Information Development Manager
- Postal service:
Oracle Corporation
Information Development Manager
500 Oracle Parkway, Mailstop 4op978
Redwood Shores, CA 94065
USA

Please indicate if you would like a reply.

If you have problems with the software, please contact your local Oracle World Wide Support Center.

[Oracle Corporation welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.](#)

[You can send comments to us in the following ways](#)

- Electronic mail — jpgcomnt@us.oracle.com
- FAX - 650-506-7225. Attn: Java Platform Group, Information Development Manager
- Postal service
Oracle Corporation
Information Development Manager
500 Oracle Parkway, Mailstop 4op978
Redwood Shores, CA 94065
USA

Please provide the following information, and indicate if you would like a reply:

Name:

Title

Company:

Department:

Electronic Mail Address:

Postal Address:

Phone Number:

Book Title:

Version Number:

If you like, you can use the following questionnaire to give us feedback:

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, please indicate the chapter, section, and page number (if available).

If you have problems with the software, please contact your local Oracle World Wide Support Center.

Preface

Who Should Read This Book

This book has been written for the following audiences:

- **Management**—You might have purchased Oracle8i for reasons other than Java development within the database. However, if you want to know more about Oracle8i Java features, see "[Overview of Oracle8i Java Documentation](#)" on page 1-22 for a management perspective.
- **Non-Java Developers**—Oracle database programming consists of PL/SQL and other non-Java programming. For experienced PL/SQL developers who are not familiar with Java, a brief overview of Java and object-oriented concepts is discussed in the first part of [Chapter 1, "Introduction to Java in Oracle8i"](#). For more detailed information on Java, see "[Java Information Resources](#)" at the end of this Preface.
- **Java Developers**—Pure Java developers are used to a Java environment that follows Sun Microsystem's specification. However, when Java is combined in the database, both Java and database concepts merge. Thus, the Java environment within Oracle8i is expanded to include database concerns. The bulk of this book discusses the differences you must understand to run Java in the database. The following outlines the two viewpoints that arise from this merge:
 - * **Java environment**—Note that Oracle8i delivers a compliant Java implementation—any 100% pure Java code will work. Oracle8i JServer affects your Java development in the way you manage your classes and the environment in which your classes exist. For example, the classes must be loaded into the database. In addition, there is a clearer separation of client and server in the Oracle8i model.

-
- * Database environment—You must be aware of database concepts for managing your Java objects. This book gives you a comprehensive view of how the two well-defined realms—Oracle8i database and Java environment—fit together. For example, when deciding on your security policies, you must consider both database security and Java security for a comprehensive security policy.

Java API Programming Models

The building blocks Java developers use in Oracle8i are as follows:

- Java stored procedures—You can develop Java applications that are stored in the database. Once loaded, these procedures can be invoked from SQL, PL/SQL, or as triggers. See the *Oracle8i Java Stored Procedures Developer's Guide* for more information.
- JDBC and SQLJ—You can write a Java application that accesses SQL data from the client or directly on the server.
- Distributed Java CORBA or EJB applications—You can develop distributed EJB or CORBA applications that are loaded and invoked in the database.

Each of these models is briefly discussed in [Chapter 1, "Introduction to Java in Oracle8i"](#) and examples are given in [Chapter 3, "Invoking Java in the Database"](#). Both of these chapters should help you decide which model to use for your particular application. Once you decide on the appropriate model, examine the appropriate developer's guide for in depth information on each model. For example, if you decide to use Java stored procedures, you should examine the book *Oracle8i Java Stored Procedures Developer's Guide*.

Java Information Resources

The following table lists the sources of current information discussed in the Java programming documentation suite:

Location	Description
http://www.oracle.com/java	The latest offerings, updates, and news for Java within the Oracle8i database. This site contains Frequently Asked Questions (FAQ), updated JDBC drivers, SQLJ reference implementations, and white papers that detail Java application development. In addition, you can download try-and-buy Java tools from this site.

Location	Description
http://java.sun.com/	Sun Microsystem's web site that is the central source for Java. This site contains Java products and information, such as tutorials, book recommendations, and the Java Developer's Kit (JDK). The JDK is located at http://java.sun.com/products
http://java.sun.com/docs/books/jls http://java.sun.com/docs/books/vmspec	The Oracle8i Java Server (JServer) is based on the Java Language (JLS) and the Java Virtual Machine (JVM) specifications.
comp.lang.java.programmer comp.lang.java.databases	Internet newsgroups can be a valuable source of information on Java from other Java developers. We recommend that you monitor these two newsgroups. Note: Oracle monitors activity on some of these newsgroups and posts responses to Oracle-specific issues.

Your local or on-line bookstore has many useful Java references. You can find another listing of materials that are helpful to beginners and that you can use as general references, in the *Oracle8i Java Stored Procedures Developer's Guide*.



Introduction to Java in Oracle8i

Java applications are supported within the Oracle8i database. Java applications can range from the simple standalone application to large, enterprise solutions using EJB or CORBA. All supported Java APIs cannot be covered within a single document; thus, several books describe the full support for Java within Oracle8i. This book provides a general overview for how you should program your Java applications when loading and running these applications in the database. Secondly, this book helps you choose which type of Java application you might develop, and direct you to the corresponding book for detailed information on that subject.

This chapter contains the following information:

- Introduces the Java language for Oracle database programmers. Oracle PL/SQL developers are accustomed to developing server-side applications that have tight integration with SQL data. You can develop Java server-side applications that take advantage of the scalability and performance of the Oracle database. If you are not familiar with Java, see "[Overview of Java](#)" on page 1-2.
- Examines why you should consider using Java within an Oracle8i database. See "[Why Use Java in Oracle8i?](#)" on page 1-12. In addition, a brief description is given of the Java application development interfaces supported within Oracle8i. These include SQLJ, JDBC, Java stored procedures, EJB, and CORBA. See "[Oracle's Java Application Strategy](#)" on page 1-17.
- Provides a roadmap to the Oracle8i Java documentation. Several Java application types are supported within Oracle8i. Each of these types are described generally in this book, and more intimately in their own books. "[Overview of Oracle8i Java Documentation](#)" on page 1-22 shows you which books cover each Java application type in detail.

Contents

- [Overview of Java](#)
- [Why Use Java in Oracle8i?](#)
- [Oracle's Java Application Strategy](#)
- [Overview of Oracle8i Java Documentation](#)

Overview of Java

Java, which was developed at Sun Microsystems, has emerged over the last several years as the object-oriented programming language of choice. It includes the following concepts:

- A Java virtual machine (JVM), which provides the fundamental basis for platform independence
- Automated storage management techniques, the most visible of which is garbage collection
- Language syntax that borrows from C and enforces strong typing

The result is a language easily learned by existing C programmers, but which remains truly object-oriented and efficient for application-level programs.

Java and Object-Oriented Programming Terminology

This section covers some basic terminology for discussing details of Java application development in the Oracle8i environment. The terms should be familiar to experienced Java programmers. A detailed discussion of object-oriented programming or of the Java language is beyond the scope of this book. Many texts, in addition to the complete language specification, are available at your bookstore and on the Internet. See "[Java Information Resources](#)" in the [Preface](#), and "Suggested Reading" in the *Oracle8i Java Stored Procedures Developer's Guide*, for pointers to reference materials and for places to find Java-related information on the Internet.

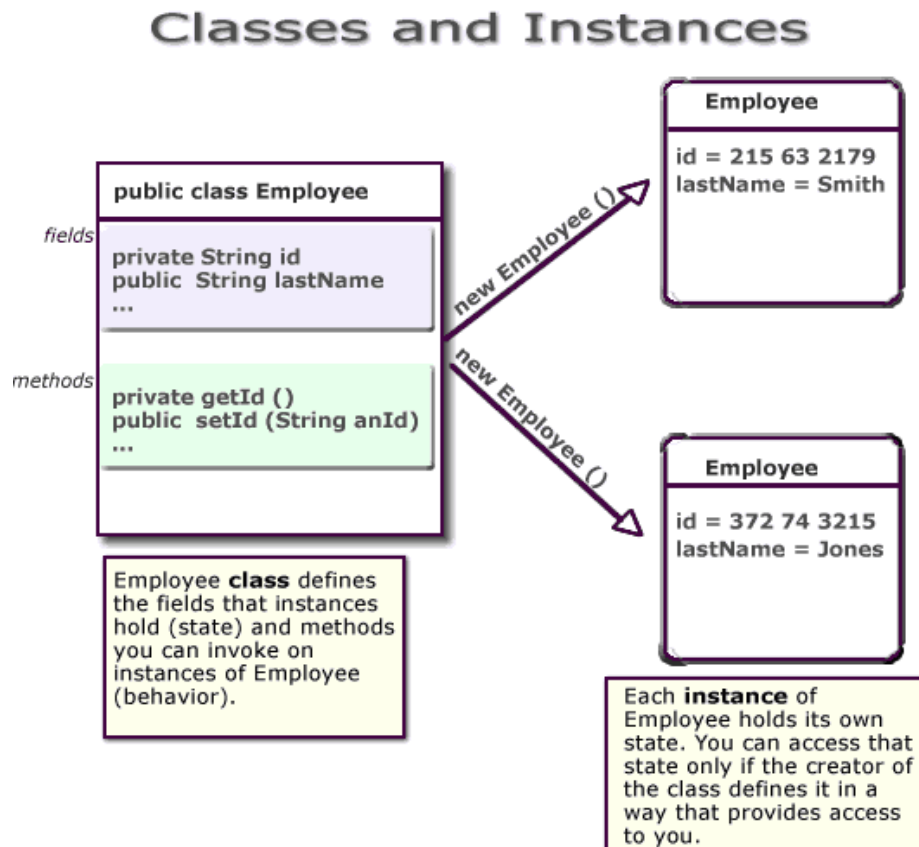
Classes

All object-oriented programming languages support the concept of a class. As with a table definition, a class provides a template for objects that share common characteristics. Each class can contain the following:

- **Attributes**—static or instance variables that each object of a particular class possesses.
- **Methods**—you can invoke methods defined by the class or inherited by any classes extended from the class.

When you create an object from a class, you are creating an instance of that class. The instance contains the fields of an object, which are known as its data, or state. [Figure 1-1](#) shows an example of an `Employee` class defined with two attributes: last name (`lastName`) and employee identifier (`ID`).

Figure 1–1 Classes and Instances



When you create an instance, the attributes store individual and private information relevant only to the employee. That is, the information contained within an employee instance is known only for that single employee. The example in [Figure 1–1](#) shows two instances of employee—Smith and Jones. Each instance contains information relevant to the individual employee.

Attributes

Attributes within an instance are known as fields. Instance fields are analogous to the fields of a relational table row. The class defines the fields, as well as the type of each field. You can declare fields in Java to be static, public, private, protected, or default access.

- Public, private, protected, or default access fields are created within each instance.
- Static fields are like global variables in that the information is available to all instances of the employee class.

The language specification defines the rules of visibility of data for all fields. Rules of visibility define under what circumstances you can access the data in these fields.

Methods

The class also defines the methods you can invoke on an instance of that class. Methods are written in Java and define the behavior of an object. This bundling of state and behavior is the essence of encapsulation, which is a feature of all object-oriented programming languages. If you define an `Employee` class, declaring that each employee's `id` is a private field, other objects can access that private field only if a method returns the field. In this example, an object could retrieve the employee's identifier by invoking the `Employee.getId()` method.

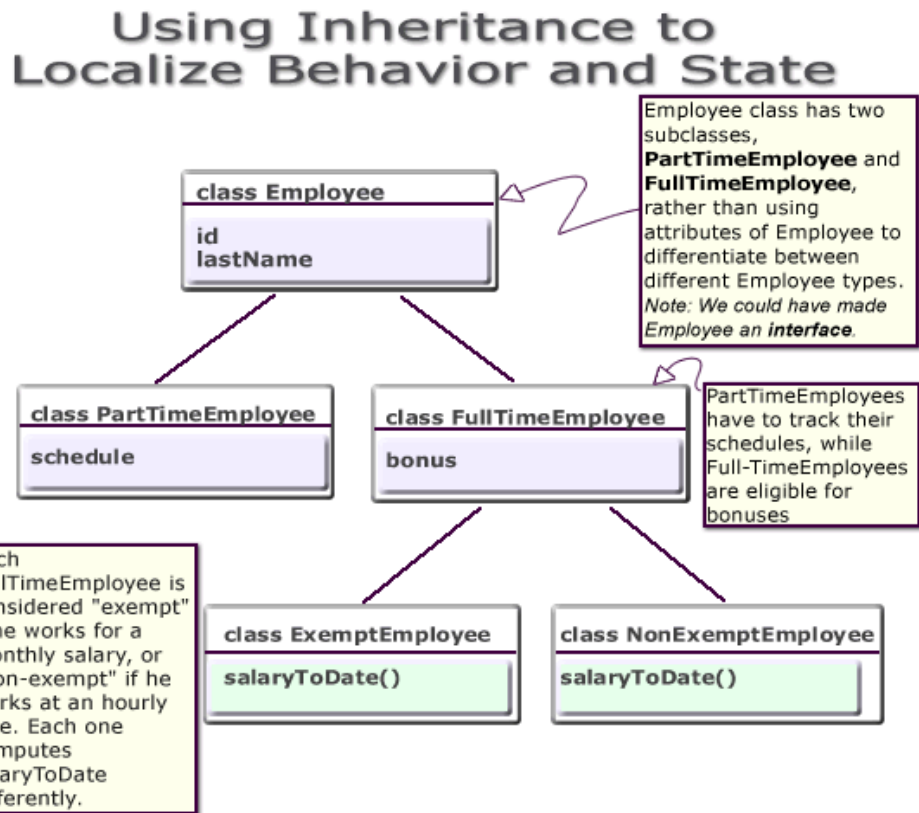
In addition, with encapsulation, you can declare that the `Employee.getId()` method is private, or you can decide not to write an `Employee.getId()` method. Encapsulation helps you write programs that are reusable and not misused. Encapsulation makes public only those features of an object that are declared public; all other fields and methods are private. Private fields and methods can be used for internal object processing.

Class Hierarchy

Java defines classes within a large hierarchy of classes. At the top of the hierarchy is the `Object` class. All classes in Java inherit from the `Object` class at some level, as you walk up through the inheritance chain of superclasses. When we say Class B inherits from Class A, each instance of Class B contains all the fields defined in class B, as well as all the fields defined in Class A. For example, in [Figure 1-2](#), the `FullTimeEmployee` class contains the `id` and `lastName` fields defined in the `Employee` class because it inherits from the `Employee` class. In addition, the `FullTimeEmployee` class adds another field, `bonus`, which is contained only within `FullTimeEmployee`.

You can invoke any method on an instance of Class B that was defined in either Class A or B. In our employee example, the `FullTimeEmployee` instance can invoke methods defined only within its own class, or methods defined within the `Employee` class.

Figure 1–2 Inheritance Hierarchy



Instances of Class B are substitutable for instances of Class A, which makes inheritance another powerful construct of object-oriented languages for improving code reuse. You can create new classes that define behavior and state where it makes sense in the hierarchy, yet make use of pre-existing functionality in class libraries.

Interfaces

Java supports only single inheritance; that is, each class has one and only one class from which it inherits. If you must inherit from more than one source, Java provides the equivalent of multiple inheritance, without the complications and confusion that usually accompany it, through interfaces. Interfaces are similar to classes; however, interfaces define method signatures, not implementations. The methods

are implemented in classes declared to implement an interface. Multiple inheritance occurs when a single class simultaneously supports many interfaces.

Polymorphism

Assume in our `Employee` example that the different types of employees must be able to respond with their compensation to date. Compensation is computed differently for different kinds of employees.

- `FullTimeEmployees` are eligible for a bonus
- `NonExemptEmployees` get overtime pay

In traditional procedural languages, you would write a long `switch` statement, with the different possible cases defined.

```
switch: (employee.type) {
    case: Employee
        return employee.salaryToDate;
    case: FullTimeEmployee
        return employee.salaryToDate + employee.bonusToDate
    ...
}
```

If you add a new kind of `Employee`, you must update your `switch` statement. If you modify your data structure, you must modify all `switch` statements that use it. In an object-oriented language such as Java, you implement a method, `compensationToDate()`, for each subclass of `Employee` class that requires any special treatment beyond what is already defined in `Employee` class. For example, you could implement the `compensationToDate()` method of `NonExemptEmployee`, as follows:

```
private float compensationToDate() {
    return super.compensationToDate() + this.overtimeToDate();
}
```

You implement `FullTimeEmployee`'s method, as follows:

```
private float compensationToDate() {
    return super.compensationToDate() + this.bonusToDate();
}
```


The common usage of the method name `compensationToDate()` allows you to invoke the identical method on different classes and receive different results, without knowing the type of employee you are using. You do not have to write a special method to handle `FullTimeEmployees` and `PartTimeEmployees`. This

ability for the different objects to respond to the identical message in different ways is known as polymorphism.

In addition, you could create an entirely new class that does not inherit from `Employee` at all—`Contractor`—and implement a `compensationToDate()` method in it. A program that calculates total payroll to date would iterate over all people on payroll, regardless of whether they were full-time, part-time, or contractors, and add up the values returned from invoking the `compensationToDate()` method on each. You can safely make changes to the individual `compensationToDate()` methods with the knowledge that callers of the methods will work correctly. For example, you can safely add new fields to existing classes.

The Java Virtual Machine (JVM)

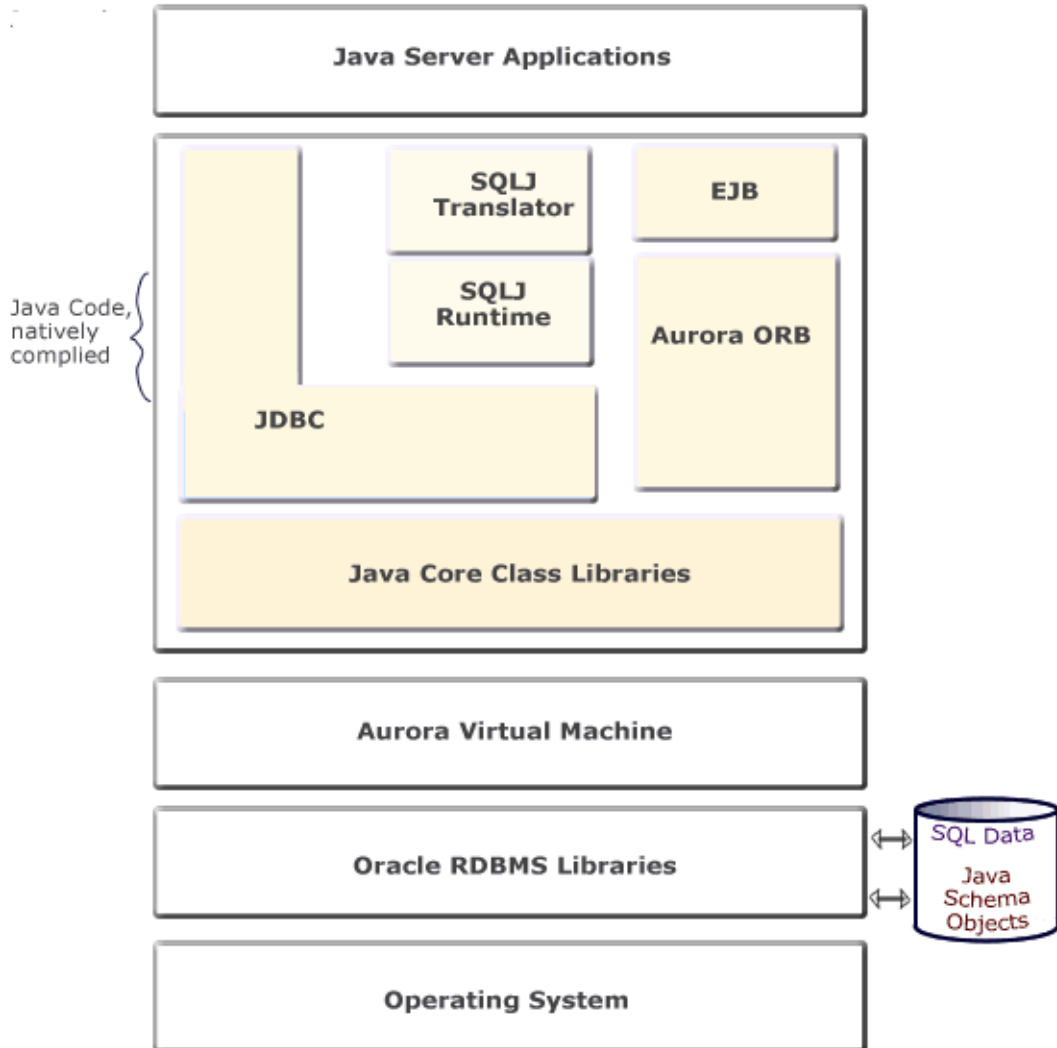
As with other high-level computer languages, your Java source compiles to low-level machine instructions. In Java, these instructions are known as bytecodes (because their size is uniformly one byte of storage). Most other languages, such as C, compile to machine-specific instructions; for example, instructions specific to an Intel or HP processor. Your Java source compiles to a standard, platform-independent set of bytecodes, which interacts with a Java virtual machine (JVM). The JVM is a separate program optimized for the specific platform on which you execute your Java code. [Figure 1-3](#) shows how Java can maintain platform independence. Your Java source is compiled into bytecodes, which are platform independent. Each platform has installed a JVM that is specific to its operating system. The Java bytecodes from your source get interpreted through the JVM into appropriate platform dependent actions.

Figure 1–3 Java Component Structure

When you develop a Java program, you use predefined core class libraries written in the Java language. The Java core class libraries are logically divided into packages that provide commonly-used functionality, such as basic language support (`java.lang`), input/output (`java.io`), and network access (`java.net`). Together, the JVM and core class libraries provide a platform on which Java programmers can develop with the confidence that any hardware and operating system that supports Java will execute their program. This concept is what drives the “write once, run anywhere” idea of Java.

[Figure 1–4](#) illustrates how Oracle’s Java applications sit on top of the Java core class libraries, which in turn sit on top of the JVM. Because Oracle’s Java support system is located within the database, the JVM interacts with the Oracle database libraries, instead of directly with the operating system.

Figure 1-4 JServer Component Structure



Sun Microsystems furnishes publicly available specifications for both the Java language and the JVM. The Java language specification (JLS) defines things such as syntax and semantics; the JVM specification defines the necessary low-level behavior for the “machine” that executes the bytecodes. In addition, Sun Microsystems provides a compatibility test suite for JVM implementors to determine if they have complied with the specifications. This test suite is known as the Java Compatibility Kit (JCK). Oracle’s JVM implementation complies fully with JCK. Part of the overall Java strategy is that an openly specified standard, together with a simple way to verify compliance with that standard, allows vendors to offer uniform support for Java across all platforms.

Key Features of the Java Language

The Java language has key features that make it ideal for developing server applications. These features include:

- **Simplicity**—Java is a simpler language to master than most others you use in server applications because of its consistent enforcement of the object model. The large, standard set of class libraries brings powerful tools to Java developers on all platforms.
- **Portability**—Java is portable across platforms. It is possible to write platform-dependent code in Java, but it is also simple to write programs that move seamlessly across machines. Oracle server applications, which do not support graphical user interfaces directly on the platform that hosts them, also tend to avoid the few platform portability issues that Java has.
- **Automatic Storage Management**—The Java virtual machine automatically performs all memory allocation and deallocation during program execution. Java programmers can neither allocate nor free memory explicitly. Instead, they depend on the JVM to perform these bookkeeping operations, allocating memory as they create new objects and deallocating memory when the objects are no longer referenced. The latter operation is known as garbage collection.
- **Strong Typing**—Before you use a Java variable, you must declare the class of the object it will hold. Java’s strong typing makes it possible to provide a reasonable and safe solution to inter-language calls in the case of Java and PL/SQL and to integrate Java and SQL.
- **No Pointers**—Although Java retains much of the flavor of C in its syntax, it does not support direct pointers or pointer manipulation. You pass all parameters, except primitive types, by reference (that is, object identity is preserved), not by value. Java does not provide C’s low level, direct access to pointers, which eliminates memory corruption and leaks.

- **Exception Handling**—Java exceptions are objects. Java requires developers to declare which exceptions can be thrown by methods in any particular class.
- **Flexible Namespace**—Java defines classes and holds them within a hierarchical structure that mirrors the Internet's domain namespace. You can distribute Java applications and avoid name collisions. Java extensions such as the Java Naming and Directory Interface (JNDI) provide a framework for multiple name services to be federated. Java's namespace approach is flexible enough for Oracle to incorporate the concept of a schema for resolving class names, while fully complying with the language specification.
- **Security**—The design of Java bytecodes and the JVM allow for built-in mechanisms to verify Java binary code has not been tampered with. Oracle8i is installed with an instance of SecurityManager, which, combined with Oracle database security, secures who can invoke any Java methods.
- **Standards for Connectivity to Relational Databases**—JDBC and SQLJ enable Java code to access and manipulate data resident in relational databases. Oracle provides drivers that allow vendor-independent, portable Java code to access the relational database.

Why Use Java in Oracle8i?

The only reason that you are allowed to write and load Java applications within the database is because it is a safe language. Java has been developed to prevent anyone tampering with the operating system that the Java code resides in. Some languages, such as C, can introduce problems within the database. Java, because of its design, is a safe language to allow within the database.

Although the Java language presents many advantages to developers, providing an implementation of a JVM that supports Java server applications in a scalable manner is a challenge. This section discusses some of these challenges.

- [Multithreading](#)
- [Automated Storage Management](#)
- [Footprint](#)
- [Performance](#)
- [Dynamic Class Loading](#)

Multithreading

Multithreading support is often cited as one of the key scalability features of the Java language. Certainly, the Java language and class libraries make it simpler to write multithreaded applications in Java than many other languages, but it is still a daunting task in any language to write reliable, scalable multithreaded code.

As a database server, Oracle8i efficiently schedules work for thousands of users. The Oracle8i Aurora JVM uses the facilities of the RDBMS server to concurrently schedule Java execution for thousands of users. Although Oracle8i supports Java language level threads required by the Java language specification (JLS) and Java Compatibility Kit (JCK), using threads within the scope of the database will not increase your scalability. Using the embedded scalability of the database eliminates the need for writing multithreaded Java servers. You should use the database's facilities for scheduling users by writing single-threaded Java applications. The database will take care of the scheduling between each application; thus, you achieve scalability without having to manage threads. You can still write multithreaded Java applications, but multiple Java threads will not increase your server's performance.

One difficulty multithreading imposes on Java is the interaction of threads and automated storage management, or garbage collection. The garbage collector executing in a generic JVM has no knowledge of which Java language threads are executing or how the underlying operating system schedules them.

- **Non-Oracle8i model**—A single user maps to a single Java language level thread; the same single garbage collector manages all garbage from all users. Different techniques typically deal with allocation and collection of objects of varying lifetimes and sizes. The result in a heavily multithreaded application is, at best, dependent upon operating system support for native threads, which can be unreliable and limited in scalability. High levels of scalability for such implementations have not been convincingly demonstrated.
- **Oracle8i JServer model**—Even when thousands of users connect to the server and execute the same Java code, each user experiences it as if he is executing his own Java code on his own Java virtual machine. The responsibility of the Oracle8i JServer is to make use of operating system processes and threads, using the scalable approach of the Oracle RDBMS. As a result of this approach, the JVM's garbage collector is more reliable and efficient because it never collects garbage from more than one user at any time. Refer to "[Threading in JServer](#)" on page 2-44 for more information on the thread model implementation in JServer.

Automated Storage Management

Garbage collection is a major feature of Java's automated storage management, eliminating the need for Java developers to allocate and free memory explicitly. Consequently, this eliminates a large source of memory leaks that commonly plague C and C++ programs. There is a price for such a benefit: garbage collection contributes to the overhead of program execution speed and footprint. Although many papers have been written qualifying and quantifying the trade-off, the overall cost is reasonable, considering the alternatives.

Garbage collection imposes a challenge to the JVM developer seeking to supply a highly scalable and fast Java platform. Aurora's JVM meets these challenges in the following ways:

- Aurora JVM uses the Oracle8i scheduling facilities, which can manage multiple users efficiently.
- Garbage collection is consistently performant for multiple users because garbage collection is focused on a single user within a single session. The Oracle8i Aurora JVM enjoys a huge advantage because the burden and complexity of the memory manager's job does not increase as the number of users increases. The memory manager performs the allocation and collection of objects within a single session—which typically translates to the activity of a single user.
- Aurora JVM uses different garbage collection techniques depending on the type of memory used. These techniques provide high efficiency and low overhead.

Footprint

The footprint of an executing Java program is affected by many factors:

- Size of the program itself—how many classes and methods and how much code they contain.
- Complexity of the program—the amount of core class libraries Aurora uses as the program executes, as opposed to the program itself.
- Amount of state Aurora uses—how many objects Aurora allocates, how large they are, and how many must be retained across calls.
- Ability of the garbage collector and memory manager to deal with the demands of the executing program, which is often non-deterministic. The speed with which objects are allocated and the way they are held on to by other objects influences the importance of this factor.

From a scalability perspective, the key to supporting many concurrent clients is a minimum per-user session footprint. Aurora keeps the per-user session footprint to a minimum by placing all read-only data for users, such as Java bytecodes, in shared memory. Appropriate garbage collection algorithms are applied against call and session memories to maintain a small footprint for the user's session. Aurora uses three types of garbage collection algorithms to maintain the user's session memory:

- Generational scavenging for short-lived objects
- Mark and lazy sweep collection for objects that exist for the life of a single call
- Copying collector for long-lived objects—objects that live across calls within a session

Performance

JServer performance is enhanced by implementing a native compiler.

How Native Compilers Improve Performance

Java executes platform-independent bytecodes on top of a JVM, which in turn deals with the specific hardware platform. Anytime you add levels within software, your performance is degraded. Because Java requires going through an intermediary to interpret platform-independent bytecodes, a degree of inefficiency exists for Java applications that does not exist within a platform-dependent language, such as C. To address this issue, several JVM suppliers create native compilers. Native compilers translate Java bytecodes into platform-dependent native code. This eliminates the interpreter step and improves performance. The following describes two methods for native compilation:

Compiler	Description
Just In Time (JIT) Compilation	JIT compilers quickly compile Java bytecodes to native (platform-specific) machine code during runtime. This does not produce an executable to be executed on the platform; instead, it provides platform-dependent code from Java bytecodes that is executed directly after it is translated. This should be used for Java code that is run frequently, which will be executed at speeds closer to languages such as C.

Compiler	Description
Static Compilation	Static compilation translates Java bytecodes to platform-independent C code before runtime. Then, a standard C compiler compiles the C code into an executable for the target platform. This approach is more suitable for Java applications that are modified infrequently. This approach takes advantage of the mature and efficient platform-specific compilation technology found in modern C compilers.

Oracle8i uses static compilation to deliver its core Java class libraries, the Aurora/ORB, and JDBC code in natively compiled form. It is applicable across all the platforms Oracle supports, whereas a JIT approach requires low-level, processor-dependent code to be written and maintained for each platform. You can use this native compilation technology with your own Java code. Refer to "[Natively Compiled Code](#)" on page 6-2 for more information.

Dynamic Class Loading

Another strong feature of Java is dynamic class loading. The class loader loads classes from the disk (and places them in the JVM-specific memory structures necessary for interpretation) only as they are used during program execution. The class loader locates the classes in the CLASSPATH and loads them during program execution. This approach, which works well for applets, poses the following problems in a server environment:

Problem	Description	Solution
Predictability	The class loading operation places a severe penalty on first-time execution. A simple program can cause Aurora to load many core classes to support its needs. A programmer cannot easily predict or determine the number of classes Aurora loads.	Aurora loads classes dynamically, just as with any other Java virtual machine. The same one-time class loading speed hit is encountered. However, because Aurora loads the classes into shared memory, no other users of those classes will cause the classes to load again—they will simply use the same pre-loaded classes.

Problem	Description	Solution
Reliability	A benefit of dynamic class loading is that it supports program updating. For example, you would update classes on a server, and clients who download the program and load it dynamically see the update whenever they next use the program. Server programs tend to emphasize reliability. As a developer, you must know that every client executes a specific program configuration. You do not want clients to inadvertently load some classes that you did not intend them to load.	Oracle8i separates the upload and resolve operation from the class loading operation at runtime. You upload Java code you developed to the server using the <code>loadjava</code> utility. Instead of using <code>CLASSPATH</code> , you specify a resolver at installation time. The resolver is analogous to <code>CLASSPATH</code> , but allows you to specify the schemas in which the classes reside. This separation of resolution from class loading means you always know what program users execute. Refer to the <i>Oracle8i Java Tools Reference</i> , for details on <code>loadjava</code> and resolvers.

Oracle's Java Application Strategy

One appeal of Java is its ubiquity and the growing number of programmers capable of developing applications using it. Oracle furnishes enterprise application developers with an end-to-end Java solution for creating, deploying, and managing Java applications. The total solution consists of client- and server-side programmatic interfaces, tools to support Java development, and a Java virtual machine integrated with the Oracle8i database server. All of these products are 100 percent compatible with Java standards.

In addition to the Aurora JVM, the Oracle8i's Java programming environment consists of:

- Java stored procedures as the Java equivalent and companion for PL/SQL. Java stored procedures are tightly integrated with PL/SQL. You can call a Java stored procedure from a PL/SQL package; you can call PL/SQL procedures from a Java stored procedure.
- SQL data can be accessed through JDBC and SQLJ programming interfaces.
- Distributed enterprise application development through an Object Request Broker (the Aurora/ORB) and Enterprise JavaBeans support.
- Dynamic HTML pages through Servlets and JavaServer Pages.
- Tools and scripts used in assisting in development, class loading, and class management.

To enable your decision making for which Java APIs to use, examine the following table:

Type of functionality you need	Java API to use
To have a Java procedure invoked from SQL, such as a trigger.	Java Stored Procedures
To invoke a static, simple SQL statement from a known table with known column names from a Java object.	SQLJ
To invoke dynamic, complex SQL statements from a Java object.	JDBC
To invoke Java from an HTML page	Servlets or JavaServer Pages
To create a multi-tier Java application.	CORBA or EJB

Java Stored Procedures

If you are a PL/SQL programmer exploring Java, you will be interested in Java stored procedures. A Java stored procedure is a program you write in Java to execute in the server, exactly as a PL/SQL stored procedure. You invoke it directly with products like SQL*Plus or indirectly with a trigger and can access it from any Net8 client—OCI, PRO*, JDBC or SQLJ. The *Oracle8i Java Stored Procedures Developer's Guide*, explains how to write stored procedures in Java, how to access them from PL/SQL, and how to access PL/SQL functionality from Java.

In addition, you can use Java to develop powerful programs independently of PL/SQL. Oracle8i provides a fully compliant implementation of the Java programming language and JVM.

PL/SQL Integration and Oracle RDBMS Functionality

You can invoke existing PL/SQL programs from Java and invoke Java programs from PL/SQL. This solution protects and leverages your existing investment while opening up the advantages and opportunities of Java-based Internet computing.

Oracle offers two different application programming interfaces (APIs) for Java developers to access SQL data—JDBC and SQLJ. Both APIs are available on client and server, so you can deploy the same code in either place.

- **JDBC Drivers**—Used to build client/server 2-tier applications.
- **SQLJ - Embedded SQL in Java**—Used to access static SQL. You must know the name of the columns.

JDBC Drivers

JDBC is a database access protocol that enables you to connect to a database and then prepare and execute SQL statements against the database. Core Java class libraries provide only one JDBC API. JDBC is designed, however, to allow vendors to supply drivers that offer the necessary specialization for a particular database. Oracle delivers JServer with the following three distinct JDBC drivers.

Driver	Description
JDBC Thin Driver	You can use the JDBC thin driver to write 100% pure Java applications and applets that access Oracle SQL data. The JDBC thin driver is especially well-suited to Web browser-based applications and applets because you can dynamically download it from a Web page just like any other Java applet.
JDBC Oracle Call Interface Driver	The JDBC Oracle Call Interface (OCI) driver accesses Oracle-specific native code (that is, non-Java) libraries on the client or middle tier, providing a richer set of functionality and some performance boost compared to the JDBC thin driver, at the cost of significantly larger size and client-side installation.
JDBC Server-side Internal Driver	Oracle8i uses the JServer server-side internal driver when Java code executes on the server. It allows Java applications executing in the server's Java virtual machine to access locally defined data (that is, on the same machine and in the same process) with JDBC. It provides a further performance boost because of its ability to use underlying Oracle RDBMS libraries directly, without the overhead of an intervening network connection between your Java code and SQL data. By supporting the same Java-SQL interface on the server, Oracle 8i does not require you to rework code when deploying it.

For more information on JDBC, see ["Utilizing SQLJ and JDBC for Querying Database"](#) on page 3-11 or a complete detailed description within the *Oracle8i JDBC Developer's Guide and Reference*.

SQLJ – Embedded SQL in Java

JDBC provides a low-level API for accessing SQL data from Java. It introduces Java classes that mirror their SQL equivalents. Oracle has worked with other vendors, including IBM, Tandem, Sybase, and Sun Microsystems, to develop a standard way to embed SQL statements in Java programs—SQLJ. This work has resulted in a new standard (ANSI x.3.135.10-1998) for a simpler and more highly productive programming API than JDBC. A user writes applications to this higher-level API and then employs a preprocessor to translate the program to standard Java source

with JDBC calls. At runtime, the program can communicate with multi-vendor databases using standard JDBC drivers. SQLJ provides a simple, but powerful, way to develop both client-side and middle-tier applications that access databases from Java. You can use it in stored procedures, triggers, methods within the JServer environment, and with EJB and CORBA. In addition, you can combine SQLJ programs with JDBC.

The SQLJ translator is a Java program that translates embedded SQL in Java source code to pure JDBC-based Java code. Because JServer provides a complete Java environment, you can not only compile SQLJ programs on a client for execution on the JServer, but you can compile them directly on the server. Oracle8i's adherence to Internet standards allows you to choose the development style that fits your needs.

For more information on SQLJ, see "[Utilizing SQLJ and JDBC for Querying Database](#)" on page 3-11 or for a complete detailed description, see the *SQLJ Developer's Guide and Reference*.

Dynamic HTML Invoking Java

You can invoke Java from within dynamic HTML protocols, such as servlets or JavaServer Pages.

Servlets

A servlet is useful for managing requests between HTML and Java applications, specifically within a web server context. Servlets are to servers what applets are to browsers. Servlets, as a true Java and HTML based protocol is platform independent, and thus can be used in multiple environments.

See the *Oracle8i Oracle Servlet Engine User's Guide* for more information.

JavaServerPages

JavaServer Pages is a method of generating dynamic content in pages that are output by a Web application, such as an application running on a Web server. This allows you to include Java code snippets and calls to external Java components within the HTML code (or other markup code, such as XML) of your Web pages. JavaServer Pages works nicely as a front-end for business logic and dynamic functionality in JavaBeans and Enterprise JavaBeans (EJBs). In addition, anything that you can include in a normal HTML page can be included in a JSP page as well.

For example, a JSP page can call a component such as a JavaBean or Enterprise JavaBean, and the bean will directly or indirectly access the database, generally through JDBC or perhaps SQLJ.

A JSP page is translated into a Java servlet before being executed. It processes HTTP requests and generates responses similarly to any other servlet. JSP pages are fully interoperable with servlets—JSP pages can include output from a servlet or forward to a servlet, and servlets can include output from a JSP page or forward to a JSP page.

See the *Oracle JavaServer Pages Developer's Guide and Reference* for more information.

Distributed Application Development

In addition to support for traditional RDBMS-stored procedures, JServer comes with a built-in CORBA 2.0 ORB and support for Enterprise JavaBeans (EJB). CORBA and EJB allow you to distribute Java components and application logic between client, middle-tier, and database server.

OMG CORBA ORB The ORB allows programs you develop in any language to communicate directly with the Oracle8i database through Internet Inter-ORB Protocol (IIOP), the standard wire-protocol defined by the Object Management Group (OMG).

Enterprise JavaBeans (EJB) For 100% pure Java applications, EJB is the standard framework for deploying component-based, secure, transactional applications on JServer.

Using EJB Components

"[Java and Object-Oriented Programming Terminology](#)" on page 1-2 discusses encapsulation as a key element of object-oriented programming. Each object maintains its own private state and supports a set of behaviors, which you implement as methods. Java provides a formal way to define components, using JavaBeans. A JavaBean component is a reusable object or group of objects (more precisely, an object graph) that you can manipulate in a builder tool of some type. IDEs, such as JDeveloper, provide tools to build user interfaces that use JavaBeans and create JavaBean components. Each bean specifies its public interface and properties that can be manipulated. JavaBeans do not have to be visually-oriented components. Virtually any Java programming abstraction can potentially be represented and manipulated as a bean.

A large component library provides the basis for assembling an application from pre-built, pre-tested building blocks. However, beans are limited in their ability to build complex business applications involving transactional logic. To address this limitation, a group of companies, including Oracle, Sun Microsystems, and IBM,

developed the Enterprise JavaBean (EJB) specification. EJB introduces a declarative mechanism for specifying how components deal with transactions and security. Refer to the *Oracle8i Enterprise JavaBeans Developer's Guide and Reference* or the *Oracle8i CORBA Developer's Guide and Reference* books for detailed information about using EJB components in Oracle8i.

There are alternative component models to JavaBeans and Enterprise JavaBeans—notably, Microsoft's COM and COM+ models. If you have existing Microsoft COM-oriented applications, they can interact with open Internet standards, such as JavaBeans and EJB, with bridge products available from different vendors.

Development Tools

The introduction of Java to the Oracle8i server allows you to use several Java Integrated Development Environments. JServer's adherence to Java compatibility and open Internet standards and protocols ensures that your 100% pure Java programs work when you deploy them on JServer. Oracle delivers JServer with many tools or utilities, all written in Java, that make development and deployment of Java server applications easier. Oracle's JDeveloper has many features designed specifically to make deployment of Java stored procedures and Enterprise JavaBeans easier.

Overview of Oracle8i Java Documentation

This guide is the starting point for Oracle8i Java developers. It outlines some of the unique features of Java programming with Oracle8i, including aspects of the Aurora JVM, explaining how to take advantage of these features in your Java programs.

Once you have mastered the basics of Java development within the Oracle8i database, you might need more information for the specific protocol you will use in implementing your Java application. The following list includes other books within the documentation set that will help you in your application development:

Protocol	Description	Book Title
JDBC	Oracle8i Java developers should become familiar with Oracle's Java Database Connectivity (JDBC) product because it provides the basis for accessing SQL data from Java programs, as well as Oracle-specific extensions to this Java standard. JDBC is an industry standard.	<i>Oracle8i JDBC Developer's Guide and Reference</i>

Protocol	Description	Book Title
SQLJ	You may find it easier to develop Java programs that access SQL data using embedded SQL in Java (SQLJ). SQLJ uses a preprocessor, written in Java, to translate embedded SQL statements to standard JDBC-style programs. SQLJ is an industry standard.	<i>Oracle8i SQLJ Developer's Guide and Reference</i>
JPublisher	JPublisher provides a simple and convenient tool to create Java programs that access existing Oracle relational database tables.	<i>Oracle8i JPublisher User's Guide</i>
Java Stored Procedures	<p>If you are a PL/SQL programmer exploring Java, you will be interested in Java stored procedures. A Java stored procedure is a program you write in Java to execute in the server, exactly as a PL/SQL stored procedure. You invoke it directly with products like SQL*Plus or indirectly with a trigger and can access it from any Net8 client—OCI, PRO*, JDBC or SQLJ. The <i>Oracle8i Java Stored Procedures Developer's Guide</i> explains how to write stored procedures in Java, how to access them from PL/SQL, and how to access PL/SQL functionality from Java.</p> <p>In addition, you can use Java to develop powerful programs independently of PL/SQL. Oracle8i provides a fully compliant implementation of the Java programming language and JVM.</p>	<i>Oracle8i Java Stored Procedures Developer's Guide.</i>
JavaServer Pages and Servlets	Dynamic HTML protocols that invoke back-end Java that exists on the database or interact with the web server,	<i>Oracle8i Oracle Servlet Engine User's Guide</i> <i>Oracle JavaServer Pages Developer's Guide and Reference</i>
EJB and CORBA	For distributed applications, you will utilize either the ORB or EJB technology. Oracle's open distributed object technology is included in its Object Request Broker (the Aurora/ORB) and Enterprise JavaBeans (EJB) functionality. The Aurora/ORB and EJB furnish powerful standards-based frameworks and tools to help you build scalable Java applications that provide seamless transactional access to Oracle data across your intranet or the Internet.	<i>Oracle8i Enterprise JavaBeans Developer's Guide and Reference</i> <i>Oracle8i CORBA Developer's Guide and Reference</i>

Writing Java Applications on Oracle8i

JServer runs standard Java applications. However, by integrating Java classes within the database server, your environment is different from a typical Java development environment. This chapter describes the basic differences for writing, installing, and deploying Java applications within Oracle8i.

- [Overview](#)
- [Database Sessions Imposed on Java Applications](#)
- [Execution Control](#)
- [Migrating from JDK 1.1 to Java 2](#)
- [Java Code, Binaries, and Resources Storage](#)
- [Preparing Java Class Methods for Execution](#)
- [User Interfaces on the Server](#)
- [Shortened Class Names](#)
- [Class.forName\(\) on JServer](#)
- [Managing Your Operating System Resources](#)
- [Threading in JServer](#)

Note: You should refer to the detailed documentation for the different JServer APIs, to fully explore their usage. The intent of this chapter is to place Java APIs in an overall context, with enough detail for you to see how they fit together and how you use them in the JServer environment.

Overview

As discussed in Chapter 1, the Oracle8i JServer platform is a standard, compatible Java environment, which will execute any 100% pure Java application. It has been implemented by Oracle to be compatible with the Java Language Specification and the Java virtual machine specification. It supports the standard Java binary format and the standard Java APIs. In addition, Oracle8i adheres to standard Java language semantics, including dynamic class loading at runtime. However, unlike other Java environments, the JServer is embedded within the Oracle8i RDBMS and, therefore, introduces a number of new concepts. This section gives an overview of the differences between Sun Microsystems's JDK environment and the environment that occurs when you combine Java within the Oracle8i database.

Terminology

Term	Definition
JServer	Java-enabled Oracle8i database server.
Aurora	Oracle8i JVM.
Session	As a user who executes Java code, you must establish a session in the server. The word <i>session</i> as we employ it here is identical to the standard Oracle (or any other database server) usage. A session is typically, although not necessarily, bounded by the time a single user connects to the server.
Call	<p>When a user causes Java code to execute within a session, we refer to it as a <i>call</i>. You can initiate a call in different ways.</p> <ul style="list-style-type: none">■ A SQL client program executes a Java stored procedure.■ A trigger can execute a Java stored procedure.■ A PL/SQL program calls some Java code.■ A CORBA client invokes a method on a CORBA object.■ An EJB client invokes a method on an EJB object. <p>In all cases, a call begins, some combination of Java, SQL, or PL/SQL code is executed to completion, and the call ends.</p>

In your standard Java environment, you run a Java application through the interpreter by executing `java <classname>`. This causes the application to execute within a process on your operating system.

With the Aurora JVM, you must load the application into the database, publish the interface, and then run the application within a database session. This book

discusses how to run your Java applications within the database. Specifically, see the following sections on instructions for Java in the database:

- Load and publish your Java applications before execution—See "[Java Code, Binaries, and Resources Storage](#)" and "[Preparing Java Class Methods for Execution](#)" starting on page 2-14.
- Running within a database session—See "[Database Sessions Imposed on Java Applications](#)" on page 2-3.

In addition, certain features, included within standard Java, change when you run your application within a database session. These are covered in the following sections:

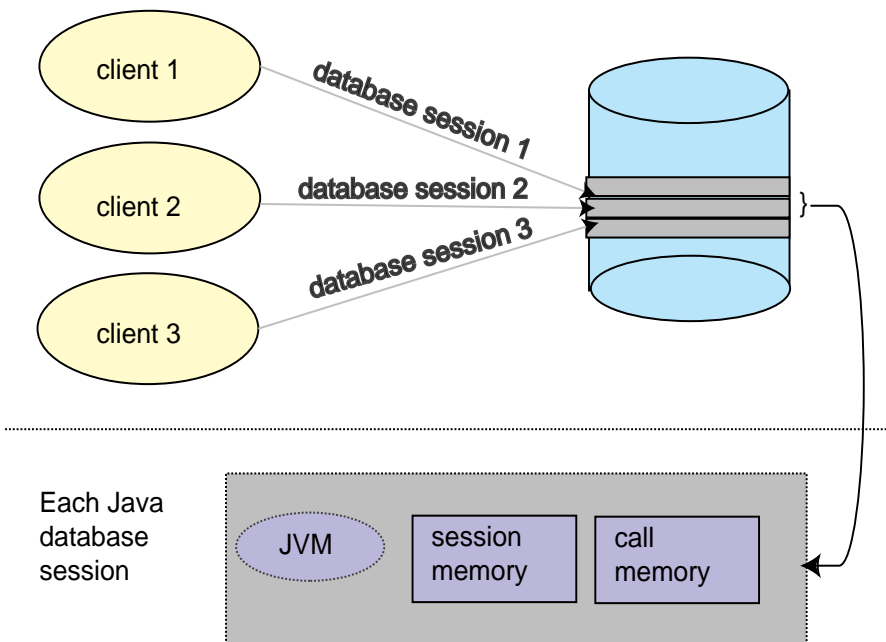
- [Execution Control](#)
- [User Interfaces on the Server](#)
- [Shortened Class Names](#)
- [Class.forName\(\) on JServer](#)
- [Managing Your Operating System Resources](#)
- [Threading in JServer](#)

Once you are familiar with this chapter, see [Chapter 3, "Invoking Java in the Database"](#) for directions on how to set up your client, and examples for invoking different types of Java applications.

Database Sessions Imposed on Java Applications

In incorporating Java within the Oracle8i database, your Java application exists within the context of a database session. JServer sessions are entirely analogous to traditional Oracle sessions. Each JServer session maintains the client's Java state across calls within the session.

As demonstrated in [Figure 2-1](#), each Java client starts up a database session as the environment for executing Java within the database. Garbage collection, session memory, and call memory exist solely for each client within its session.

Figure 2–1 Java environment within each database session

Within the context of a session, the client performs the following:

1. Connects to the database and opens a session.
2. Executes Java within the database. This is referred to as a call.
3. Continues to work within the session performing as many calls as necessary.
4. Ends the session.

Within a single session, the client has its own Java environment, which is separate from every other client's environment. It appears to the client as if a separate, individual JVM was invoked for each session, although the implementation is vastly more efficient than this seems to imply. Within a session, the Aurora JVM manages the scalability for you within the database. Every call executed from a single client is managed within its own session—separately from other clients. The Aurora JVM maximizes sharing read-only data between clients and emphasizes a minimum amount of per-session incremental footprint to maximize performance for multiple clients.

The underlying server environment hides the details associated with session, network, state, and other shared resource management issues from Java server code. Static variables are all local to the client. No client can access another client's static variables, because the memory is not available across session boundaries. Each client executes its calls within its own session, so each client's activities are separate from any other client. During a call, you can store objects in static fields of different classes, and you can expect this state to be available for your next call. The entire state of your Java program is private to you and exists for your entire session.

The Aurora JVM manages the following within the session:

- all the objects referenced by Java static variables, all the objects referred to by these objects, and so on (their transitive closure)
- garbage collection for the single client
- session memory for static variables and across call memory needs
- call memory for variables that exist within a single call

Session Lifetime

When you connect to Oracle8i, you start a database session. A session ends when one of the following events occurs:

1. The user invokes the `oracle.aurora.mts.session.Session.THIS_SESSION().endSession()` method.
2. The session times out. This is optional for CORBA or EJB sessions.
3. The user takes some action outside of Java code to end the database session.

Java Supported APIs

For the current Oracle8i release, we offer five Java APIs—Java stored procedures, Servlets, JavaServer Pages, CORBA distributed objects, and Enterprise JavaBeans (EJBs).

API	Lifetime
Java stored procedures	The lifetime of a Java stored procedure session is identical to the SQL session in which it is embedded. This concept is familiar to PL/SQL users. Any state represented in Java transparently persists for the lifetime of the RDBMS session, simplifying the process of writing stored procedures, triggers, and methods for Oracle Abstract Data Types. Individual invocations of Java code within a session are known as calls. For example, a call may be initiated by a SQL call.
Servlets and JavaServer Pages	Servlets and JavaServer pages provide a dynamic method for invoking Java from within HTML. Used to manage communication between HTML and Java applications stored on the resource, whether middle-tier cache or database.
CORBA and EJB	CORBA and EJB provide a more object-oriented style of message sending between clients and servers. Clients must implicitly or explicitly establish a session in the server. Every message you send on the client to a server-resident object initiates a call. Refer to the <i>Oracle8i Enterprise JavaBeans Developer's Guide and Reference</i> or the <i>Oracle8i CORBA Developer's Guide and Reference</i> books for specifics.

Note: The concepts of call and session apply across all uses of Oracle8i.

In addition, you can access SQL data through SQLJ or JDBC. See [Chapter 3, "Invoking Java in the Database"](#) for examples of each Java API.

Execution Control

In Sun Microsystems's JDK environment, you develop Java applications with a `main()` method, which is called by the interpreter when the class is run. The `main()` method is invoked when you execute `java <classname>` on the command-line. This command starts the java interpreter and passes the desired classname to be executed to the interpreter. The java interpreter loads the class and starts the execution by invoking `main()`. However, Java applications within the database do not start their execution from a `main()` method.

After loading your Java application within the database (see "[Loading Classes](#)" on page 2-23), you can execute your Java code by invoking any static method within the loaded class. The class or methods must be published for you to execute them

(see "Publishing" on page 2-30). Your only entry point is no longer always assumed to be `main()`. Instead, when you execute your Java application, you specify a method name within the loaded class as your entry point.

For example, in a normal Java environment, you would start up the Java object on the server by executing the following:

```
java myprogram
```

where `myprogram` is the name of a class that contains a `main()` method. In `myprogram`, `main()` immediately calls `mymethod` for processing incoming information.

In Oracle8i, you load the `myprogram.class` file into the database and publish `mymethod` as an entry-point. Then, the client or trigger can invoke `mymethod` explicitly.

Migrating from JDK 1.1 to Java 2

Java 2 is, for the most part, compatible with JDK 1.1. Sun Microsystems changed certain features, such as the security feature, in Java 2. These changes are documented at the following Sun Microsystems's web site:

<http://java.sun.com/products/jdk/1.2/compatibility.html>

The following sections discuss how the changes made within Java 2 affected Oracle8i:

- [Your Development Environment](#)
- [JDBC 2.0](#)
- [Java 2 Security](#)
- [Java 2 ORB APIs](#)

Your Development Environment

The level of your development environment determines your interoperability with the server. If your development environment is Java 2-based, any code compiled and debugged on your system can be loaded and executed on the database. However, if you are developing applications in a JDK 1.1 development environment on your client, you can only use JDK 1.1 classes. With very few exceptions, you can load and execute your JDK 1.1 application in the Java 2 database server. Of course,

you always have the option to code your Java 2-based application on your system, load it into the database, and use the Java 2 compiler that exists on the database.

Note: There is another workaround for using Java 2 security even though your code is JDK 1.1-based. The security APIs are provided within a PL/SQL package. You can use these call specifications before your code executes; thus, enabling the correct Java 2 permissions.

JDBC 2.0

Even though 8.1.5 was JDK 1.1-based, JDBC 2.0 support was added to 8.1.5 in an Oracle-specific package—`oracle.jdbc2`. However, the current version of Oracle8i supports Java2, so JDBC 2.0 exists in its intended package—`java.sql`. If you have JDBC programs that used the `oracle.jdbc2` package for JDBC 2.0 APIs, you must modify these programs before executing with JDBC 2.0 drivers.

Note: If your Java application does not use the `oracle.jdbc2` package JDBC APIs (or return an object whose type is within this package), no migration is necessary to connect to the 8.1.6 database.

With the addition of Java 2 in Oracle8i, the JDBC 2.0 support exists in the `java.sql` package, which is contained in the JDK core libraries. This affects your application, as follows:

Environment

Running a server application in Oracle8i which uses JDBC 2.0.	You must use the <code>java.sql</code> package. If you use <code>oracle.jdbc2</code> , your application will not compile.
Running a client in a Java 2 environment	You must use the <code>java.sql</code> package. If you use <code>oracle.jdbc2</code> , your client will not compile.

Environment

Running a client in a JDK 1.1 environment.

You can continue to use the `oracle.jdbc2` APIs, but only if you compile against `classes111.zip` and you use this code only for accessing the server as a client. If you try to load this code within the server, you must migrate to the `java.sql` package.

Server Applications Using JDBC 2.0

Since the server application is loaded and runs within the database, which is always a Java 2 environment, you must port your application, as stated below:

1. Replace all imports and other mentions of `oracle.jdbc2` with `java.sql` in your programs.

The `oracle.jdbc2` package contains the JDBC 2.0 implementation that Oracle implemented in the JDK 1.1 drivers. Because those classes and interfaces are available in Java 2, `oracle.jdbc2` is not included in `classes12.zip`.

2. The return type of the `getTypeMap()` method of `Connection` has been changed from `java.util.Dictionary` to `java.util.Map`. Modify your application accordingly. No change is necessary if you are using `java.util.Hashtable`, because `Hashtable` implements `java.util.Map`.
3. Replace `classes111.zip` with `classes12.zip` in your makefile.
4. Recompile and relink your executable.

You must recompile all classes that used to import `oracle.jdbc2`. In addition, you must recompile any classes where an object is returned and its type exists within the `oracle.jdbc2` package. Consider the following example:

```
#import java.sql.*;
. . .
jd = (java.sql.Connection) getDC();
```

If the `getDC` method returns an object type from within the `oracle.jdbc2` package, the compiler will not catch this as a problem because it is typecast to `java.sql.*`. Instead, you will receive the following error when you load the application in the database:

```
ORA-29521: referenced name oracle/jdbc2/<classname> could not be found.
```

Recompile all of these classes against `classes12.zip`, which will correct the object type returned to `java.sql.*`.

For more information, See Chapter 4 in the *Oracle8i JDBC Developer's Guide and Reference*.

Clients Using JDBC 2.0

The following client applications can interoperate with 8.1.6:

- Java client conforms to JDK 1.1 and uses the JDBC 2.0 APIs contained in `oracle.jdbc2`. This application imports the `oracle.jdbc2` package within `classes111.zip`.
- Java client conforms to Java 2 and uses the JDBC 2.0 APIs contained in `java.sql`. This application imports the `java.sql` package contained within the `classes12.zip`, not the `oracle.jdbc2` package.

Java 2 Security

Java 2 security is implemented in 8.1.6. The JDK 1.1 security sandbox is no longer applicable within Oracle8i. To use the Java 2 security Permissions without modifying your code, you can manage these Permissions through the PL/SQL package—`DBMS_JAVA`. To execute any of the Java 2 security methods, such as `doPrivileged`, you must compile your application within a Java 2 environment, as the Permissions are new for the JDK 1.2 release.

See "[Java 2 Security](#)" on page 5-3 for more information on Java 2 security.

Java 2 ORB APIs

Oracle8i JServer updated its ORB implementation to Visigenic 3.4. This version is compatible with both JDK 1.1 and Java 2.

Note: All existing CORBA applications must regenerate their stubs and skeletons to work with 8.1.6. You must use the 8.1.6 tools when regenerating your application.

Sun Microsystems's Java 2 contains an ORB implementation; JDK 1.1 did not. Thus, when you imported the Visigenic libraries and invoked the CORBA methods, it always invoked the Visigenic implementation. With Java 2, if you invoke the

CORBA methods without any modifications—as discussed below—you will invoke Sun Microsystems’s CORBA implementation, which can cause unexpected results.

The following lists the three methods for accessing CORBA server objects in Oracle8i from your client and the recommendations for bypassing Sun Microsystems’s CORBA implementation:

- **JNDI Lookup**—The setup for the lookup method is identical for both JDK 1.1 and Java 2. However, you must regenerate the stubs and skeletons.
- **Aurora ORB Interface**—The Aurora ORB provides an interface for initializing the ORB. If you do not use JNDI, your client initializes an ORB on its node to communicate with the ORB in the database. You can use an Aurora ORB on your client through this class.
- **CORBA ORB Interface**—If you want to use OMG’s CORBA ORB interface, you must set a few properties to ensure you are accessing the correct implementation. If you do not wish to use the Aurora ORB on your client, you can use the pure CORBA interfaces. However, you need to set up your environment to direct your calls to the correct implementation.

JNDI Lookup

If you are using JNDI on the client to access CORBA objects that reside in the server, no code changes are required. However, you must regenerate your CORBA stubs and skeletons.

Aurora ORB Interface

If your client environment uses JDK 1.1, you do not need to change your existing code. You will need to regenerate your stubs and skeletons.

If your client environment has been upgraded to Java 2, you can initialize the ORB through the `oracle.aurora.jndi.orb_dep.Orb.init` method. This method guarantees that when you initialize the ORB, it will initialize only a single ORB instance. That is, if you use the Java 2 ORB interface, it returns you a new ORB instance each time you invoke the `init` method. Aurora’s `init` method initializes a singleton ORB instance. Each successive call to `init` returns an object reference to the existing ORB instance.

In addition, the Aurora ORB interface manages the session-based IIOP connection.

oracle.aurora.jndi.orb_dep.Orb Class There are several `init` methods, each with a different parameter list. The following describes the syntax and parameters for each `init` method.

Note: The returned class for each `init` method are different. You can safely cast the `org.omg.CORBA.ORB` class to `com.visigenic.vbroker.orb.ORB`.

```
public com.visigenic.vbroker.orb.ORB init();
public org.omg.CORBA.ORB init(Properties props);
public org.omg.CORBA.ORB init(String[] args, Properties props);
```

Parameter	Description
Properties props	ORB system properties.
String[] args	Arguments that are passed to the ORB instance.

Example 2-1 Using Aurora ORB init method

The following example shows a client instantiating an ORB using the Aurora Orb class.

```
// Create the client object and publish it to the orb in the client
// Substitute Aurora's Orb.init for OMG ORB.init call
// old way: org.omg.CORBA.ORB orb = org.omg.CORBA.ORB.init ();
com.visigenic.vbroker.orb.ORB orb = oracle.aurora.jndi.orb_dep.Orb.init();
```

Joining a Session If your client receives a reference to an object that is created in a session, it can invoke a method on that object within the session. However, since all clients must authenticate, you must provide a username and password to the database. If the server requires client-side authentication in the form of `SSL_CREDENTIALS`, you can provide the client's username, password, and role, which is passed on the connect handshake within the `ORB.init` method.

```
public org.omg.CORBA.ORB init(String un, String pw, String role,
                             boolean ssl, java.util.Properties props);
```

Parameter	Description
String un	The username for client-side authentication.
String pw	The password for client-side authentication.
String role	The role to use after logging on.
Boolean ssl	If true, SSL is enabled for the connection. If false, a NON-SSL connection is used.

Parameter	Description
Properties props	ORB system properties.

CORBA ORB Interface

If you have implemented a pure CORBA client—that is, you do not use JNDI—you need to set the following properties before the ORB initialization call. These properties direct the call to the Aurora implementation, rather than the Java 2 implementation. This ensures the behavior that you expect. The behavior expected from the Visigenic ORB is as follows:

- Even if you invoke `ORB.init` more than once, only a single ORB instance is created. If you do not set these properties, be aware that each invocation of `ORB.init` will create a new ORB instance.
- The session IIOP connection is managed correctly.
- Callbacks from the server are managed correctly.

Note: The Aurora CORBA implementation is based upon Visibroker 3.4.

Property	Assign Value
<code>org.omg.CORBA.ORBClass</code>	<code>com.visigenic.vbroker.orb.ORB</code>
<code>org.omg.CORBA.ORBSingletonClass</code>	<code>com.visigenic.vbroker.orb.ORB</code>

Example 2–2 Assigning Visigenic values to OMG properties

The following example shows how to set up the OMG properties for directing the OMG CORBA `init` method to the Visigenic implementation.

```
System.getProperties().put("org.omg.CORBA.ORBClass",
    "com.visigenic.vbroker.orb.ORB");
System.getProperties().put("org.omg.CORBA.ORBSingletonClass",
    "com.visigenic.vbroker.orb.ORB");
```

Or you can set the properties on the command line, as follows:

```
java -Dorg.omg.CORBA.ORBClass=com.visigenic.vbroker.orb.ORB
    -Dorg.omg.CORBA.ORBSingletonClass=com.visigenic.vbroker.orb.ORB
```

Backwards Compatibility for 8.1.5 CORBA and EJB Applications

The tools provided with Oracle8i, such as `publish`, have been modified to work with either a JDK 1.1 or Java 2 environment. However, any CORBA or EJB code that has been generated or loaded with the 8.1.5 version of any tool, will not succeed. Make sure that you always use the 8.1.6 version of all tools. This rule applies to your CORBA stubs and skeletons. You must regenerate all stubs and skeletons with the 8.1.6 IDL compiler.

Java Code, Binaries, and Resources Storage

In the Sun Microsystems Java development environment, Java source code, binaries, and resources are stored as files in a file system.

- Source code files are known as `.java` files.
- Compiled Java binary files are known as `.class` files.
- Resources are any data files, such as `.properties` or `.ser` files held within the file system hierarchy, which are loaded or used at runtime.

In addition, when you execute Java, you specify a `CLASSPATH`, which is a set of a file system tree roots containing your files. Java also provides a way to group these files into a single archive form—a ZIP or JAR file.

Both of these concepts are different within the database. The following describes how JServer handles Java classes and locates dependent classes:

Java code, binaries, and resources	In the JServer environment, source, classes, and resources reside within the Oracle8i database. Because they reside in the database, they are known as Java schema objects, where a schema corresponds to a database user. There are three types of Java objects: source, class, and resource. There are no <code>.java</code> , <code>.class</code> , <code>.sqlj</code> , <code>.properties</code> , or <code>.ser</code> files on the server; instead these files map to source, class, and resource Java schema objects.
------------------------------------	--

Locating Java classes	Instead of a <code>CLASSPATH</code> , you use a resolver to specify one or more schemas to search for source, class, and resource Java schema objects.
-----------------------	--

The terms `call` and `session`, used during our discussions, are not Java terms; but are server terms that apply to the Oracle8i JServer platform. The Aurora memory manager preserves Java program state throughout your session (that is, between calls). The JServer uses the Oracle database to hold Java source, classes, and

resources within a schema—Java schema objects. You can use a resolver to specify how Java, when executed in the server, locates source code, classes, and resources.

Preparing Java Class Methods for Execution

For your Java methods to be executed, you must do the following:

1. Decide when your source is going to be compiled.
2. Decide if you are going to use the default resolver or another resolver for locating supporting Java classes within the database.
3. Load the classes into the database. If you do not wish to use the default resolver for your classes, you should specify a separate resolver on the load command.
4. Publish your class or method.

Compiling Java Classes

Compilation of your source can be performed in one of the following ways:

- You can compile the source explicitly on your client machine before loading it into the database through a Java compiler, such as `javac`.
- You can ask the database to compile the source during the loading process managed within the `loadjava` tool.
- You can force the compilation to occur dynamically at runtime.

Note: If you decide to compile through `loadjava`, you can specify compiler options. See "[Specifying Compiler Options](#)" on page 2-16 for more information.

Compiling Source through `javac`

You can compile your Java with a conventional Java compiler, such as `javac`. After compilation, you load the compiled binary into the database, rather than the source itself. This is a better option, because it is normally easier to debug your Java code on your own system, rather than debugging it on the database.

Compiling Source through `loadjava`

When you specify the `-resolve` option on `loadjava` for a source file, the following occurs:

1. The source file is loaded as a source schema object.
2. The source file is compiled.
3. Class schema objects are created for each class defined in the compiled `.java` file.
4. The compiled code is stored in the class schema objects.

JServer logs all compilation errors both to `loadjava`'s logfile and the `USER_ERRORS` view. For more information on the `USER_ERRORS` view, see the *Oracle8i Reference* for a description of this table.

Compiling Source at Runtime

When you load the Java source into the database without the `-resolve` option, JServer compiles the source automatically when the class is needed during runtime. The source file is loaded into a source schema object.

JServer logs all compilation errors both to `loadjava`'s logfile and the `USER_ERRORS` view. For more information on the `USER_ERRORS` view, see the *Oracle8i Reference* for a description of this table.

Specifying Compiler Options

There are two ways to specify options to the compiler.

- Specify compiler options on the `loadjava` command line. You can specify the encoding option on the `loadjava` command line.
- Specify persistent compiler options in a per-schema database table called `JAVA$OPTIONS`. Every time you compile, the compiler uses these options. However, any specified compiler options on the `loadjava` command override the options defined in this table.

You must create this table yourself if you wish to specify compiler options this way. See "[Compiler Options Specified in a Database Table](#)" on page 2-17 for instructions to create the `JAVA$OPTIONS` table.

The following sections describe your compiler options:

- [Default Compiler Options](#)
- [Compiler Options on the Command Line](#)
- [Compiler Options Specified in a Database Table](#)

Default Compiler Options When compiling a source schema object for which there is neither a `JAVA$OPTIONS` entry nor a command line value for an option, the compiler assumes a default value as follows:

- `encoding = latin1`
- `online = true`: See the *Oracle8i SQLJ Developer's Guide and Reference* for a description of this option, which applies only to Java sources that contain SQL constructs.
- `debug = true`: This option is equivalent to `javac -g`.

Compiler Options on the Command Line The `loadjava` compiler option, `encoding`, identifies the encoding of the `.java` file. This option overrides any matching value in the `JAVA$OPTIONS` table. The values are identical to the `javac -encoding` option. This option is relevant only when loading a source file.

Compiler Options Specified in a Database Table Each `JAVA$OPTIONS` row contains the names of source schema objects to which an option setting applies; you can use multiple rows to set the options differently for different source schema objects.

You can set `JAVA$OPTIONS` entries by means of the following functions and procedures, which are defined in the database package `DBMS_JAVA`:

- `PROCEDURE set_compiler_option(name VARCHAR2, option VARCHAR2, value VARCHAR2);`
- `FUNCTION get_compiler_option(name VARCHAR2, option VARCHAR2) RETURNS VARCHAR2;`
- `PROCEDURE reset_compiler_option(name VARCHAR2, option VARCHAR2);`

The parameters for these methods are described below:

Parameter	Description
name	The name parameter is a Java package name, a fully qualified class name, or the empty string. When the compiler searches the <code>JAVA\$OPTIONS</code> table for the options to use for compiling a Java source schema object, it uses the row whose name most closely matches the schema object's fully qualified class name. A name whose value is the empty string matches any schema object name.
option	The option parameter is either 'online', 'encoding' or 'debug'. For the values you can specify for these options, see the <i>Oracle8i SQLJ Developer's Guide and Reference</i> .

A schema does not initially have a `JAVA$OPTIONS` table. To create a `JAVA$OPTIONS` table, use the `DBMS_JAVA` package's `java.set_compiler_option` procedure to set a value. The procedure will create the table if it does not exist. Specify parameters in single quotes. For example:

```
SQL> execute dbms_java.set_compiler_option('x.y', 'online', 'false');
```

Table 2–1 represents a hypothetical `JAVA$OPTIONS` database table. The pattern match rule is to match as much of the schema name against the table entry as possible. The schema name with a higher resolution for the pattern match is the entry that applies. Because the table has no entry for the `encoding` option, the compiler uses the default or the value specified on the command line. The `online` option shown in the table matches schema object names as follows:

- The name `a.b.c.d` matches class and package names beginning with `a.b.c.d`; the packages and classes are compiled with `online = true`.
- The name `a.b` matches class and package names beginning with `a.b`. The name `a.b` does not match `a.b.c.d`; therefore, the packages and classes are compiled with `online = false`.
- All other packages and classes match the empty string entry and are compiled with `online = true`.

Table 2–1 Example JAVA\$OPTIONS Table

JAVA\$OPTIONS Entries			Match Examples
Name	Option	Value	
<code>a.b.c.d</code>	<code>online</code>	<code>true</code>	<ul style="list-style-type: none"> ■ <code>a.b.c.d</code>—matches the pattern exactly. ■ <code>a.b.c.d.e</code>—first part matches the pattern exactly; no other rule matches full name.
<code>a.b</code>	<code>online</code>	<code>false</code>	<ul style="list-style-type: none"> ■ <code>a.b</code>—matches the pattern exactly ■ <code>a.b.c.x</code>—first part matches the pattern exactly; no other rule matches beyond specified rule name.
(empty string)	<code>online</code>	<code>true</code>	<ul style="list-style-type: none"> ■ <code>a.c</code>—no pattern match with any defined name; defaults to (empty string) rule ■ <code>x.y</code>—no pattern match with any defined name; defaults to (empty string) rule

Automatic Recompilation

JServer provides a dependency management and automatic build facility that will transparently recompile source programs when you make changes to the source or binary programs upon which they depend. Consider the following cases:

```
public class A
{
    B b;
    public void assignB () {b = new B()}
}
public class B
{
    C c;
    public void assignC () {c = new C()}
}
public class C
{
    A a;
    public void assignA () {a = new A()}
}
```

The system tracks dependencies at a class level of granularity. In the preceding example, you can see that classes A, B, and C depend on one another, because A holds an instance of B, B holds an instance of C, and C holds an instance of A. If you change the definition of class A by adding a new field to it, the dependency mechanism in Oracle8i flags classes B and C as invalid. Before you use any of these classes again, Oracle8i attempts to resolve them again and recompile, if necessary. Note that classes can be recompiled only if source is present on the server.

The dependency system enables you to rely on Oracle8i to manage dependencies between classes, to recompile, and to resolve automatically. You must only force compilation and resolution yourself only if you are developing and you want to find problems early. The `loadjava` utility also provides the facilities for forcing compilation and resolution if you do not want to allow the dependency management facilities to perform this for you.

Resolving Class Dependencies

Many Java classes contain references to other classes, which is the essence of reusing code. A conventional Java virtual machine searches for classes, ZIP, and JAR files within the directories specified in the CLASSPATH. In contrast, the Aurora Java virtual machine searches database schemas for class objects. With Oracle8i, you

load all Java classes within the database, so you might need to specify where to find the dependent classes for your Java class within the database.

All classes loaded within the database are referred to as class schema objects and are loaded within certain schemas. All JVM classes, such as `java.lang.*`, are loaded within `PUBLIC`. If your classes depend upon other classes you have defined, you will probably load them all within your own schema. For example, if your schema is `SCOTT`, the database resolver (the database replacement for `CLASSPATH`) searches the `SCOTT` schema before `PUBLIC`. The listing of schemas to search is known as a resolver spec. Resolver specs are per-class, whereas in a classic Java virtual machine, `CLASSPATH` is global to all classes.

When locating and resolving the interclass dependencies for classes, the resolver marks each class as valid or invalid, depending on if all interdependent classes are located or not. If the class that you load contains a reference to a class that is not found within the appropriate schemas, the class is listed as invalid. Unsuccessful resolution at runtime produces a “class not found” exception. Furthermore, runtime resolution can fail for lack of database resources if the tree of classes is very large.

Note: As with the Java compiler, `loadjava` resolves references to classes, but not to resources. Be sure to correctly load the resource files your classes need.

For each interclass reference in a class, the resolver searches the schemas specified by the resolver spec for a valid class schema object that satisfies the reference. If all references are resolved, the resolver marks the class valid. A class that has never been resolved, or has been resolved unsuccessfully, is marked invalid. A class that depends on a schema object that becomes invalid is also marked invalid.

To make searching for dependent classes easier, Oracle8i provides a default resolver and resolver spec that searches first the definer’s schema and then `PUBLIC`. This covers most of the classes loaded within the database. However, if you are accessing classes within a schema other than your own or `PUBLIC`, you must define your own resolver spec.

- loading using Oracle’s default resolver, which searches the definer’s schema and `PUBLIC`:

```
loadjava -resolve
```

- loading using your own resolver spec definition containing the `SCOTT` schema, `OTHER` schema, and `PUBLIC`:

```
loadjava -resolve -resolver "((* SCOTT)(* OTHER)(* PUBLIC))"
```

The `-resolver` option specifies the objects to search within the schemas defined. In the example above, all class schema objects are searched within `SCOTT`, `OTHER`, and `PUBLIC`. However, if you wanted to search for only a certain class or group of classes within the schema, you could narrow the scope for the search. For example, to search only for the classes `"my/gui/*"` within the `OTHER` schema, you would define the resolver spec as follows:

```
loadjava -resolve -resolver '((* SCOTT) ("my/gui/*" OTHER) (* PUBLIC))'
```

The first parameter within the resolver spec is for the class schema object; the second parameter defines the schema to search for these class schema objects.

Allowing References to Non-Existent Classes

You can specify a special option within a resolver spec that allows an unresolved reference to a non-existent class. Sometimes, internal classes are never used within a product. For example, some ISVs do not remove all references to internal test classes from the JAR file before shipping. In a normal Java environment, this is not a problem, because as long as the methods are not called, Sun Microsystems's JVM ignores them. However, the Oracle8i resolver tries to resolve all classes referenced within the JAR file—even unused classes. If the reference cannot be validated, the classes within the JAR file are marked as invalid.

To ignore references, you can specify the `"-"` wildcard within the resolver spec. The following example specifies that any references to classes within `"my/gui"` are to be allowed even if it is not present within the resolver spec schema list.

```
loadjava -resolve -resolver '((* SCOTT) (* PUBLIC) ("my/gui/*" -))'
```

In addition, you can define that all classes not found are to be ignored. Without the wildcard, if a dependent class is not found within one of the schemas, your class is listed as invalid and cannot be run. However, this is also dangerous, because if there is a dependent class on a used class, you mark a class as valid that can never run without the dependent class. In this case, you will receive an exception at runtime.

To ignore all classes not found within `SCOTT` or `PUBLIC`, specify the following resolver spec:

```
loadjava -resolve -resolver "((* SCOTT) (* PUBLIC) (* -))"
```

Note: Never use a resolver containing “-” if you later intend to load the classes that were causing you to use such a resolver in the first place. Instead, include all referenced classes in the schema before resolving.

ByteCode Verifier

According to the JVM specification, `.class` files are subject to verification before the class they define is available in a JVM. In JServer, the verification process occurs at class resolution. The resolver might find one of the following problems and issue the appropriate Oracle error code:

- ORA-29545 If the resolver determines that the class is malformed, the resolver does not mark it valid. When the resolver rejects a class, it issues an ORA-29545 error (badly formed class). The `loadjava` tool reports the error. For example, this error is thrown if the contents of a `.class` file are not the result of a Java compilation or if the file has been corrupted.
- ORA-29552 In some situations, the resolver allows a class to be marked valid, but will replace bytecodes in the class to throw an exception at runtime. In these cases, the resolver issues an ORA-29552 (verification warning), which `loadjava` will report. The `loadjava` tool issues this warning when the Java Language Specification would require an `IncompatibleClassChangeError` be thrown. JServer relies on the resolver to detect these situations, supporting the proper runtime behavior the JLS requires.

The resolver also issues warnings, as defined below:

- Resolvers containing “-”

This type of resolver marks your class valid regardless of whether classes it references are present. Because of inheritance and interfaces, you might want to write valid Java methods that use an instance of a class as if it were an instance of a superclass or of a specific interface. When the method being verified uses a reference to class A as if it were a reference to class B, the resolver must check that A either extends or implements B. For example, consider the potentially valid method below, whose signature implies a return of an instance of B, but whose body returns an instance of A:

```
B myMethod(A a) { return a; }
```

The method is valid only if A extends B or A implements the interface B. If A or B have been resolved using a “-” term, the resolver does not know that this method is safe. It will replace the bytecodes of `myMethod` with bytecodes that throw an Exception if `myMethod` is ever called.

- Use of other resolvers

The resolver ensures that the class definitions of A and B are found and resolved properly if they are present in the schemas they specifically identify. The only time you might consider using the alternative resolver is if you must load an existing JAR file containing classes that reference other non-system classes not included in the JAR file.

For more information on class resolution and loading your classes within the database, see the *Oracle8i Java Tools Reference*.

Loading Classes

This section gives an overview of the main points you should understand when loading your classes into the database. It discusses various options for the `loadjava` tool, but does not go into all the details. You can also execute `loadjava` within your SQL. See the *Oracle8i Java Tools Reference* for complete information on `loadjava`.

Unlike a conventional Java virtual machine, which compiles and loads from files, the Aurora Java virtual machine compiles and loads from database schema objects.

<code>.java</code> source files or <code>.sqlj</code> source files	correspond to Java source schema objects
<code>.class</code> compiled Java files	correspond to Java class schema objects
<code>.properties</code> Java resource files, <code>.ser</code> SQLJ profile files, or data files	correspond to Java resource schema objects

You must load all classes or resources into the database to be used by other classes within the database. In addition, at loadtime, you define who can execute your classes within the database.

The `loadjava` tool performs the following for each type of file:

Schema object	loadjava operations on object
<code>.java</code> source files	<ol style="list-style-type: none"> 1. It creates a source schema object within the definer's schema unless another schema is specified 2. It loads the contents of the source file into a schema object 3. It creates a class schema objects for all classes defined in the source file 4. If <code>-resolve</code> is requested, it does the following: <ol style="list-style-type: none"> a. It compiles the source schema object b. It resolves the class and its dependencies c. It stores the compiled class into a class schema object
<code>.sqlj</code> source files	<ol style="list-style-type: none"> 1. It creates a source schema object within the definer's schema unless another schema is specified 2. It loads contents of the source file into the schema object 3. It creates a class schema objects for all classes and resources defined in the source file 4. If <code>-resolve</code> is requested, it does the following: <ol style="list-style-type: none"> a. It translates and compile the source schema object b. It stores the compiled class into a class schema object c. It stores profile into <code>.ser</code> resource schema object and customizes it
<code>.class</code> compiled Java files	<ol style="list-style-type: none"> 1. It creates a class schema object within the definer's schema unless another schema is specified 2. It loads the class file into the schema object 3. It resolves and verify the class and its dependencies if <code>-resolve</code> is specified

<code>.properties</code> Java resource files	<ol style="list-style-type: none"> 1. It creates a resource schema object within the definer's schema unless another schema is specified 2. It loads resource file into a schema object
<code>.ser</code> SQLJ profile	<ol style="list-style-type: none"> 1. It creates a resource schema object within the definer's schema unless another schema is specified 2. It loads <code>.ser</code> resource file into a schema object and customizes it

The `dropjava` tool performs the reverse of the `loadjava` tool: it deletes schema objects that correspond to Java files. You should always use `dropjava` to delete a Java schema object created with `loadjava`. Dropping with SQL DDL commands will not update auxiliary data maintained by `loadjava` and `dropjava`. You can also execute `dropjava` from within SQL commands.

Note: More options for `loadjava` are available. However, this section discusses only the major options. See the *Oracle8i Java Tools Reference* for complete information on `loadjava` and `dropjava`.

You must abide by certain rules when loading classes into the database, which are detailed in the following sections:

- [Two Definitions of the Same Class](#)
- [Need Database Privileges and JVM Permissions](#)
- [Loading JAR or ZIP Files](#)

After loading, you can access the `USER_OBJECTS` view in your database schema to verify that your classes and resources loaded properly. For more information, see "[Checking Java Uploads](#)" on page 2-28.

Two Definitions of the Same Class

You cannot have two different definitions for the same class. This rule affects you in two ways:

- You can load either a particular Java `.class` file or its `.java` file, but not both. JServer tracks whether you loaded a class file or a source file. If you wish to update the class, you must load the same type of file that you originally loaded.

If you wish to update the other type, you must drop the first before loading the second. For example, if you loaded `x.java` as the source for class `y`, to load `x.class`, you must first drop `x.java`.

- You cannot define the same class within two different schema objects within the same schema. For example, suppose `x.java` defines class `y` and you want to move the definition of `y` to `z.java`. If `x.java` has already been loaded, `loadjava` rejects any attempt to load `z.java` (which also defines `y`). Instead, do either of the following:
 - Drop `x.java`, load `z.java` (which defines `y`), then load the new `x.java` (which does not define `y`).
 - Load the new `x.java` (which does not define `y`), then load `z.java` (which defines `y`).

Need Database Privileges and JVM Permissions

You must have the following SQL database privileges to load classes:

- `CREATE PROCEDURE` and `CREATE TABLE` privileges to load into your schema.
- `CREATE ANY PROCEDURE` and `CREATE ANY TABLE` privileges to load into another schema.
- `oracle.aurora.security.JServerPermission.loadLibraryInClass.<classname>`. See ["Permission for Loading Classes"](#) on page 5-25 for more information.

Loading JAR or ZIP Files

The `loadjava` tool accepts `.class`, `.java`, `.properties`, `.sqlj`, `.ser`, `.jar`, or `.zip` files. The JAR or ZIP files can contain source, class, and data files. When you pass `loadjava` a JAR or ZIP file, `loadjava` opens the archive and loads its members individually. There is no JAR or ZIP schema object. If the JAR or ZIP content has not changed since the last time it was loaded, it is not reloaded; therefore, there is little performance penalty for loading JAR or ZIP files. In fact, loading JAR or ZIP files is the simplest way to use `loadjava`.

Note: JServer does not reload a class if it has not changed since the last load. However, you can force a class to be reloaded through the `loadjava -force` option.

How to Grant Execute Rights

When you are loading your classes, you can grant execution rights to another user through an option for `loadjava`. There are two methods for defining who can execute your class:

- granting execution rights to a certain user or schema

The classes that define a Java application are stored within the Oracle8i RDBMS under the SQL schema of their owner. By default, classes that reside in one user's schema are not executable by other users, because of security concerns. You can allow other users (schemas) the right to execute your class through the `loadjava -grant` option.
- specifying what user or schema is evaluated for the execution rights to the class

You can specify that the user evaluated for execution rights is either the invoker or the definer. Invoker's and definer's rights are SQL concepts that are applicable to your Java classes. In SQL, when you execute under definer's rights, you execute with the user identifier for the schema that loaded the PL/SQL. With invoker's rights, your SQL executes with the user identifier for the schema that actually invokes the PL/SQL. This applies to the Java classes in that the schema that is evaluated for JVM security permission to execute is based on whether the class was loaded under invoker's or definer's rights.

Figure 2–2 Invoker’s Versus Definer’s Rights



Method invocation: Class A invokes class B; class B invokes class C.

Execution rights for classes:

- * Class A has execution rights for B.
- * Class A does not have execution rights for C.
- * Class B has execution rights for C.

With the example in [Figure 2–2](#), which class, A or B, is checked when the method in class C is executed? This depends on invoker’s or definer’s rights.

- * **Invoker’s rights**—By default, classes run under the effective identity and rights of the client’s schema that initially invokes the method in C. In this case, the rights for class A is checked. Because class A does not have permission to execute C, the request is rejected. This is the default.
- * **Definer’s rights**—Classes execute under the effective identity of the user that loaded the class. If, when you loaded class B, you specified definer’s rights, class B is checked when the method in class C is invoked. Because class B has the right to execute methods in class C, the method completes successfully. Class B is loaded with definer’s rights by executing `loadjava -definer`.

For information on JVM security permissions, see [Chapter 6, "Oracle8i Java Application Performance"](#).

Checking Java Uploads

You can query the database view `USER_OBJECTS` to obtain information about schema objects—including Java sources, classes, and resources—that you own. This allows you, for example, to verify that sources, classes, or resources that you load are properly stored into schema objects.

Columns in `USER_OBJECTS` include those contained in [Table 2-2](#) below.

Table 2-2 Key `USER_OBJECT` Columns

Name	Description
<code>OBJECT_NAME</code>	name of the object
<code>OBJECT_TYPE</code>	type of the object (such as <code>JAVA_SOURCE</code> , <code>JAVA_CLASS</code> , or <code>JAVA_RESOURCE</code>)
<code>STATUS</code>	status of the object (<code>VALID</code> or <code>INVALID</code>) (always <code>VALID</code> for <code>JAVA_RESOURCE</code>)

Object Name and Type

An `OBJECT_NAME` in `USER_OBJECTS` is the short name. The full name is stored as a short name if it exceeds 31 characters. See ["Shortened Class Names"](#) on page 2-32 for more information on full and short names.

If the server uses a short name for a schema object, you can use the `LONGNAME ()` routine of the server `DBMS_JAVA` package to receive it from a query in full name format, without having to know the short name format or the conversion rules.

```
SQL*Plus> SELECT dbms_java.longname(object_name) FROM user_objects
           WHERE object_type='JAVA_SOURCE';
```

This routine shows you the Java source schema objects in full name format. Where no short name is used, no conversion occurs, because the short name and full name are identical.

You can use the `SHORTNAME ()` routine of the `DBMS_JAVA` package to use a full name as a query criterion, without having to know whether it was converted to a short name in the database.

```
SQL*Plus> SELECT object_type FROM user_objects
           WHERE object_name=dbms_java.shortname('known_fullname');
```

This routine shows you the `OBJECT_TYPE` of the schema object of the specified full name. This presumes that the full name is representable in the database character set.

```
SVRMGR> select * from javasnm;
SHORT                                LONGNAME
-----
/78e6d350_BinaryExceptionHandl sun/tools/java/BinaryExceptionHandler
/b6c774bb_ClassDeclaration      sun/tools/java/ClassDeclaration
/af5a8ef3_JarVerifierStream1     sun/tools/jar/JarVerifierStream$1
```

Status

STATUS is a character string that indicates the validity of a Java schema object. A source schema object is VALID if it compiled successfully; a class schema object is VALID if it was resolved successfully. A resource schema object is always VALID, because resources are not resolved.

Example: Accessing USER_OBJECTS The following SQL*Plus script accesses the USER_OBJECTS view to display information about uploaded Java sources, classes, and resources.

```
COL object_name format a30
COL object_type format a15
SELECT object_name, object_type, status
       FROM user_objects
       WHERE object_type IN ('JAVA SOURCE', 'JAVA CLASS', 'JAVA RESOURCE')
       ORDER BY object_type, object_name;
```

You can optionally use wildcards in querying USER_OBJECTS, as in the following example.

```
SELECT object_name, object_type, status
       FROM user_objects
       WHERE object_name LIKE '%Alerter';
```

This routine finds any OBJECT_NAME entries that end with the characters: Alerter.

For more information about USER_OBJECTS, see the *Oracle8i Java Stored Procedures Developer's Guide*.

Publishing

Oracle8i enables clients and SQL to invoke Java methods loaded within the database, once published. You publish either the object itself or individual methods, depending on the type of Java application it is, as shown below:

Java API	Publishing method	Reference
Java stored procedures	If you write a Java stored procedure that you intend to invoke with a trigger, directly or indirectly in SQL DML or in PL/SQL, you must publish individual methods within the class. You specify how to access it through a call specification. Java programs consist of many methods in many classes; however, only a few static methods are typically exposed with call specifications.	<i>Oracle8i Java Stored Procedures Developer's Guide.</i>
Servlet and JavaServer Pages	Publish the servlet and JavaServer Pages URL within the JNDI namespace.	<i>Oracle8i Oracle Servlet Engine User's Guide</i> <i>Oracle JavaServer Pages Developer's Guide and Reference</i>
CORBA and EJB development	<p>You do not use call specifications for CORBA or EJB objects. Instead, you publish the object reference for the client to retrieve. Once the object is retrieved, the client can invoke specific methods within the object.</p> <p>Oracle8i's CORBA and EJB implementations support standard CORBA and Java styles of exposing objects by name, with accompanying CORBA and Java-style specifications of the interfaces to those objects.</p> <ul style="list-style-type: none"> ■ You publish CORBA IOR's through the <code>publish</code> tool. ■ You publish EJB Home and Remote interfaces through the <code>deployejb</code> tool. 	<i>Oracle8i Enterprise JavaBeans Developer's Guide and Reference</i> or the <i>Oracle8i CORBA Developer's Guide and Reference</i> books

User Interfaces on the Server

Oracle8i furnishes all core Java class libraries on the server, including those associated with presentation of user interfaces (`java.awt` and `java.applet`). It is, however, inappropriate for code executing in the server to attempt to bring up or materialize a user interface in the server. Imagine thousands of users worldwide exercising an Internet application that executes code that requires someone to click on a dialog presented on the server hardware. You can write Java programs that reference and use `java.awt` classes as long as you do not attempt to materialize a user interface.

When building applets, you test them using the `java.awt` and the `Peer` implementation, which is a platform-specific set of classes for support of a specific windowing system. When the user downloads an applet, it dynamically loads the proper client `Peer` libraries, and the user sees a display appropriate for the operating system or windowing system in use on the client side. Oracle8i takes the

same approach. We provide an Oracle-specific `Peer` implementation that throws an exception, `oracle.aurora.awt.UnsupportedOperation`, if you execute Java code on the Oracle8i server that attempts to materialize a user interface.

Oracle8i's lack of support for materializing user interfaces in the server means that we do not pass the Java 2 Compatibility Kit tests for `java.awt`, `java.awt.manual`, and `java.applet`. In the Oracle RDBMS, all user interfaces are supported only on client applications, although they might be displayed on the same physical hardware that supports the server—for example, in the case of Windows NT. Because it does not make sense for the server to support user interfaces, we exclude these tests from our complete Java Compatibility Kit testing.

A similar issue exists for vendors of Java-powered embedded devices and in handheld devices (known as Personal Java). Future releases of Java and the Java Compatibility Kit will provide improved factorization of user interface support so that vendors of Java server platforms can better address this issue.

Shortened Class Names

Each Java source, class, and resource is stored in its own schema object in the server. The name of the schema object is derived from the fully qualified name, which includes relevant path or package information. Dots are replaced by slashes. These fully qualified names (with slashes)—used for loaded sources, loaded classes, loaded resources, generated classes, and generated resources—are referred to in this chapter as schema object *full names*.

Schema object names, however, have a maximum of only 31 characters, and all characters must be legal and convertible to characters in the database character set. If any full name is longer than 31 characters or contains illegal or non-convertible characters, the Oracle8i server converts the full name to a *short name* to employ as the name of the schema object, keeping track of both names and how to convert between them. If the full name is 31 characters or less and has no illegal or inconvertible characters, then the full name is used as the schema object name.

Because Java classes and methods can have names exceeding the maximum SQL identifier length, Oracle8i uses abbreviated names internally for SQL access. Oracle8i provides a method within the `DBMS_JAVA` package for retrieving the original Java class name for any truncated name.

```
FUNCTION longname (shortname VARCHAR2) RETURN VARCHAR2
```

This function returns the longname from a Java schema object. An example is to print the fully qualified name of classes that are invalid for some reason.

```
select dbms_java.longname (object_name) from user_objects
       where object_type = 'JAVA CLASS' and status = 'INVALID';
```

In addition, you can specify a full name to the database by using the `shortname()` routine of the `DBMS_JAVA` package, which takes a full name as input and returns the corresponding short name. This is useful when verifying that your classes loaded by querying the `USER_OBJECTS` view.

```
FUNCTION shortname (longname VARCHAR2) RETURN VARCHAR2
```

Refer to the *Oracle8i Java Stored Procedures Developer's Guide* for a detailed example of the use of this function and ways to determine which Java schema objects are present on the server.

Class.forName() on JServer

The Java Language Specification provides the following description of `Class.forName()`:

Given the fully-qualified name of a class, this method attempts to locate, load, and link the class. If it succeeds, a reference to the `Class` object for the class is returned. If it fails, a `ClassNotFoundException` is thrown.

Class lookup is always on behalf of a referencing class through a `ClassLoader`. The difference between the JDK implementation and JServer's implementation is the method on which the class is found:

- The JDK uses one `ClassLoader` that searches the set of directory tree roots specified by the environment variable `CLASSPATH`.
- JServer defines several resolvers, which define how to locate classes. Every class has a resolver associated with it, and each class can, potentially, have a different resolver. When you execute a method that calls `Class.forName()`, the resolver of the currently executing class (`this`) is used to locate the class. See ["Resolving Class Dependencies"](#) on page 2-19 for more information on resolvers.

You can receive unexpected results if you try to locate a class with an unexpected resolver. For example, if a class `X` in schema `X` requests a class `Y` in schema `Y` to look up class `Z`, you can experience an error if you expected class `X`'s resolver to be used. Because class `Y` is performing the lookup, the resolver associated with class `Y` is used to locate class `Z`. In summary, if the class exists in another schema and you specified different resolvers for different classes—as would happen by default if they are in different schemas—you might not find the class.

You can solve this resolver problem as follows:

- Avoid any class name lookup by passing the `Class` object itself.
- Supply the `ClassLoader` in the `Class.forName` method.
- Supply the class and the schema it resides in to JServer's `classForNameAndSchema` method.
- Supply the schema and class name to `Class.forName.lookupClass`.
- Serialize your objects with the schema name with the class name.

Note: Another unexpected behavior can occur if system classes invoke `Class.forName()`. The desired class is only found if it resides in `SYS` or in `PUBLIC`. If your class does not exist in either `SYS` or `PUBLIC`, you can declare a `PUBLIC` synonym for the class.

Supply the `ClassLoader` in `Class.forName`

JServer uses resolvers for locating classes within schemas. Every class has a specified resolver associated with it. Each class can have a different resolver associated with it. Thus, the locating of classes is dependent on the definition of the associated resolver. The `ClassLoader` knows which resolver to use based upon the class specified. When you supply a `ClassLoader` to `Class.forName()`, your class is looked up in the schemas defined within the resolver of the class. The syntax for this variant of `Class.forName` is as follows:

```
Class.forName (String name, boolean initialize, ClassLoader loader);
```

The following examples show how to supply the class loader of either the current class instance or the calling class instance.

Example 2-3 Retrieve Resolver from Current Class

You can retrieve the class loader of any instance through the `Class.getClassLoader` method. The following example retrieves the class loader of the class represented by instance `x`.

```
Class c1 = Class.forName (x.whatClass(), true, x.getClass().getClassLoader());
```


Example 2-4 Retrieve Resolver from Calling Class

You can retrieve the class of the instance that invoked the executing method through the `oracle.aurora.vm.OracleRuntime.getCallerClass` method. Once you retrieve the class, invoke the `Class.getClassLoader` method on the returned class. The following example retrieves the class of the instance that invoked the `workForCaller` method. Then, its class loader is retrieved and supplied to the `Class.forName` method. Thus, the resolver used for looking up the class is the resolver of the calling class.

```
void workForCaller() {
    ClassLoader c1 =
        oracle.aurora.vm.OracleRuntime.getCallerClass().getClassLoader();
    ...
    Class c = Class.forName (name, true, c1);
}
```

Supply Class and Schema Names to classForNameAndSchema

You can resolve the problem of where to find the class by either supplying the resolver, which knows the schemas to search, or by supplying the schema in which the class is loaded. If you know in which schema the class is loaded, you can use the `classForNameAndSchema` method. JServer provides a method in the `DbmsJava` class, which takes in both the name of the class and the schema that the class resides in. This method locates the class within the designated schema.

Example 2-5 Providing Schema and Class Names

The following example shows how you can save the schema and class names in the `save` method. Both names are retrieved and the class is located using the `DbmsJava.classForNameAndSchema` method.

```
import oracle.aurora.rdbms.ClassHandle;
import oracle.aurora.rdbms.Schema;
import oracle.aurora.rdbms.DbmsJava;

void save (Class c1) {
    ClassHandle handle = ClassHandle.lookup(c1);
    Schema schema = handle.schema();
    writeName (schema.getName());
    writeName (c1.getName());
}
```

```
Class restore() {
    String schemaName = readName();
    String className = readName();
    return DbmsJava.class.forNameAndSchema (schemaName, className);
}
```

Supply Class and Schema Names to lookupClass

You can supply a single String, containing both the schema and class names, to the `oracle.aurora.util.Class.forName.lookupClass` method. When invoked, this method locates the class in the specified schema. The string must be in the following format:

```
"<schema>:<class>"
```

For example, to locate `com.package.myclass` in schema `SCOTT`, you would execute the following:

```
oracle.aurora.util.Class.forName.lookupClass("SCOTT:com.package.myclass");
```

Note: You must uppercase the schema name. In this case, the schema name is case-sensitive.

Supply Class and Schema Names when Serializing

When you de-serialize a class, part of the operation is to lookup a class based on a name. In order to ensure that the lookup is successful, the serialized object must contain both the class and schema names.

JServer provides the following classes for serializing and deserializing objects:

- `oracle.aurora.rdbms.DbmsObjectOutputStream`

This class extends `java.io.ObjectOutputStream` and adds schema names in the appropriate places.

- `oracle.aurora.rdbms.DbmsObjectInputStream`

This class extends `java.io.ObjectInputStream` and reads streams written by `DbmsObjectOutputStream`. This class can be used on any environment. If used within JServer, the schema names are read out and used when performing the class lookup. If used on a client, the schema names are ignored.

Class.forName Example

The following example shows several methods for looking up a class.

- To use the resolver of this instance's class, invoke `lookupWithClassLoader`. This method supplies a class loader to the `Class.forName` method in the `from` variable. The class loader specified in the `from` variable defaults to this class.
- To use the resolver from a specific class, call `ForName` with the designated class name followed by `lookupWithClassLoader`. The `ForName` method sets the `from` variable to the specified class. The `lookupWithClassLoader` method uses the class loader from the specified class.
- To use the resolver from the calling class, invoke the `ForName` method without any parameters. It sets the `from` variable to the calling class. Then, invoke the `lookupWithClassLoader` to locate the class using the resolver of the calling class.
- To lookup a class in a specified schema, invoke the `lookupWithSchema` method. This provides the class and schema name to the `classForNameAndSchema` method.

```
import oracle.aurora.vm.OracleRuntime;
import oracle.aurora.rdbms.Schema;
import oracle.aurora.rdbms.DbmsJava;

public class ForName {

    private Class from;
    /* Supply an explicit class to the constructor */
    public ForName(Class from) {
        this.from = from;
    }
    /* Use the class of the code containing the "new ForName()" */
    public ForName() {
        from = OracleRuntime.getCallerClass();
    }

    /* lookup relative to Class supplied to constructor */
    public Class lookupWithClassLoader(String name) throws ClassNotFoundException
    {
        /* A ClassLoader uses the resolver associated with the class*/
        return Class.forName(name, true, from.getClassLoader());
    }
}
```

```
/* In case the schema containing the class is known */
static Class lookupWithSchema(String name, String schema) {
    Schema s = Schema.lookup(schema);
    return DbmsJava.class.forNameAndSchema(name, s);
}
}
```

Managing Your Operating System Resources

Operating system resources are a limited commodity on any computer. Because Java is targeted at providing a computing platform as well as a programming language, it contains platform-independent classes and frameworks for accessing platform-specific resources. The Java class methods access operating system resources through the JVM. Java has potential problems with this model, because programmers rely on the garbage collector to manage all resources, when all that the garbage collector manages is Java objects, not the operating system resources that the Java object holds on to.

In addition, because the Aurora JVM is embedded in the database, your operating system resources, which are contained within Java objects, can be invalidated if they are maintained across calls within a session.

The following sections discuss these potential problems:

- [Overview of Operating System Resources](#)
- [Garbage Collection and Operating System Resources](#)
- [Operating System Resources Affected Across Calls](#)

Overview of Operating System Resources

In general, your operating system resources contain the following:

- | | |
|--------|---|
| memory | Aurora manages memory internally, allocating memory as you create new objects and freeing objects as you no longer need them. The language and class libraries do not support a direct means to allocate and free memory. " Automated Storage Management " on page 1-14 discusses garbage collection. |
| files | Java contains classes that represent file resources. Instances of these classes hold on to your operating system's file constructs, such as file handles, which can become invalid between calls in a session. |

sockets	Java contains classes that represent socket resources. Instances of these classes hold on to socket constructs, some of which can become invalid between calls in a session. See " Sockets " on page 2-43 for information specific to maintaining sockets across calls.
threads	Threads are discouraged within the Aurora JVM because of scalability issues. However, you can have a multi-threaded application within the database. " Threading in JServer " on page 2-44 discusses in detail the Aurora's JVM threading model.

Operating System Resource Access

By default, a Java user does not have direct access to most operating system resources. A system administrator may give permission to a user to access these resources by modifying the JVM security restrictions. The JVM security enforced upon system resources conforms to Java 2 security. See "[Java 2 Security](#)" on page 5-3 for more information.

Operating System Resource Lifetime

You access operating system resources using the standard core Java classes and methods. Once you access a resource, the time that it remains active (usable) varies according to the type of resource.

Resource	Lifetime
Files	The system closes all files left open when a database call ends.
Memory	Memory is garbage collected as described in " Automated Storage Management " on page 1-14.
Threads	All threads are terminated when a call ends.
Objects that depend on operating system resources	<p>Regardless of the usable lifetime of the object (for example, the defined lifetime for a thread object), the Java object can be valid for the duration of the session. This can occur, for example, if the Java object is stored in a static class variable or a class variable references it directly or indirectly. If you attempt to use one of these Java objects after its usable lifetime is over, Aurora throws an exception. This is true for the following examples:</p> <ul style="list-style-type: none"> ■ If an attempt is made to read from a <code>java.io.FileInputStream</code> that was closed at the end of a previous call, a <code>java.io.IOException</code> is thrown. ■ <code>java.lang.Thread.isAlive()</code> is false for any <code>Thread</code> object running in a previous call and still accessible in a subsequent call.

Resource	Lifetime
Sockets	<ul style="list-style-type: none"> ■ Sockets can exist across calls. ■ ServerSockets on an MTS server terminate when the call ends. ■ ServerSockets on a dedicated server can exist across calls. <p>See "Sockets" on page 2-43 more information.</p>

Garbage Collection and Operating System Resources

Imagine that memory is divided up into two realms: Java object memory and operating system constructs. The Java object memory realm contains all objects and variables. Operating system constructs include resources that the operating system allocates to the object when it asks. These resources include files, sockets, and so on.

Basic programming rules dictate that you close all memory—both Java objects and operating system constructs. Java programmers incorrectly assume that all memory is freed by the garbage collector. The garbage collector was created to collect all unused Java object memory. However, it does not close any operating system constructs. All operating system constructs must be closed by the program before the Java object is collected.

For example, whenever an object opens a file, the operating system creates the file and gives the object a file handle. If the file is not closed, the operating system will hold the file handle construct open until the call ends or JVM exits. This can cause you to run out of these constructs earlier than necessary. There are a finite number of handles within each operating system. To guarantee that you do not run out of handles, close your resources before exiting the method. This includes closing the streams attached to your sockets. You should close the streams attached to the socket before closing the socket.

So why not expand the garbage collector to close all operating system constructs? For performance reasons, the garbage collector cannot examine each object to see if it contains a handle. Thus, the garbage collector collects Java objects and variables, but does not issue the appropriate operating system methods for freeing any handles.

[Example 2-6](#) shows how you should close the operating system constructs.

Example 2-6 Closing your operating system resources

```
public static void addFile(String[] newFile) {
    File inFile = new File(newFile);
    FileReader in = new FileReader(inFile);
    int i;

    while ((i = in.read()) != -1)
        out.write(i);
    /*closing the file, which frees up the operating system file handle*/
    in.close();
}
```

If you do not close the `in` file, eventually the `File` object will be garbage collected. However, even if the `File` object is garbage collected, the operating system still believes that the file is in use, because it was not closed.

Note: You might want to use Java finalizers to close resources. However, finalizers are not guaranteed to run in a timely manner. Instead, finalizers are put on a queue to execute when the garbage collector has time. If you close your resources within your finalizer, it might not be freed up until the JVM exits. The best approach is to close your resources within the method.

Operating System Resources Affected Across Calls

You should close resources that are local to a single call when the call ends. However, for static objects that hold on to operating system resources, you must be aware of how these resources are affected after the call ends.

The JVM automatically closes any open operating system constructs—in [Example 2-7](#), the file handle—when the call ends. This can affect any operating system resources within your Java object. For example, if you have a file opened within a static variable, the file handle is closed at the end of the call for you. So, if you hold on to the `File` object across calls, the next usage of the file handle throws an exception.

In [Example 2-7](#), class `Concat` enables multiple files to be written into a single file, `outFile`. On the first call, `outFile` is created. The first input file is opened, read, input into `outFile`, and the call ends. Because `outFile` is statically defined, it is moved into session space between call invocations. However, the file handle—that

is, the `FileDescriptor`—is closed at the end of the call. The next time you call `addFile`, you will get an exception.

Example 2–7 Compromising your operating system resources

```
public class Concat {
    static File outFile = new File("outme.txt");
    FileWriter out = new FileWriter(outFile);

    public static void addFile(String[] newFile) {
        File inFile = new File(newFile);
        FileReader in = new FileReader(inFile);
        int i;

        while ((i = in.read()) != -1)
            out.write(i);
        in.close();
    }
}
```

There is a workaround. To make sure that your handles stay valid, you should close your files, buffers, and so on, at the end of every call; reopen the resource at the beginning of the next call. Another option is to use the database, rather than using operating system resources. For example, try to use database tables, rather than a file. Or simply do not store operating system resources within static objects expected to live across calls; use operating system resources only within objects local to the call.

[Example 2–8](#) shows how you can perform concatenation, as in [Example 2–7](#), without compromising your operating system resources. The `addFile` method opens the `outme.txt` file within each call, making sure that anything written into the file is appended to the end. At the end of each call, the file is closed. Two things occur:

1. The `File` object no longer exists outside of a call.
2. The operating system resource, the `outme.txt` file, is reopened for each call. If you had made the `File` object a static variable, the closing of `outme.txt` within each call would ensure that the operating system resource is not compromised.

Example 2-8 Correctly managing your operating system resources

```

public class Concat {

public static void addFile(String[] newFile) {
    /*open the output file each call; make sure the input*/
    /*file is written out to the end by making it "append=true"*/
    FileWriter out = new FileWriter("outme.txt", TRUE);
    File inFile = new File(newFile);
    FileReader in = new FileReader(inFile);
    int i;

    while ((i = in.read()) != -1)
        out.write(i);
    in.close();
    /*close the output file between calls*/
    out.close();
    }
}

```

Sockets

Sockets are used in setting up a connection between a client and a server. For each database connection, sockets are used at either end of the connection. Your application does not set up the connection; the connection is set up by the underlying networking protocol: Net8's TTC or IIOP. See "[Configuring JServer](#)" on page 4-6 for information on how to configure your connection.

You might also wish to set up another connection—for example, connecting to a specified URL from within one of the classes stored within the database. To do so, instantiate sockets for servicing the client and server sides of the connection.

- The `java.net.Socket()` constructor creates a client socket.
- The `java.net.ServerSocket()` constructor creates a server socket.

A socket exists at each end of the connection. The server-side of the connection that listens for incoming calls is serviced by a `ServerSocket`. The client-side of the connection that sends requests is serviced through a `Socket`. You can use sockets as defined within the JVM with the following restriction: a `ServerSocket` instance within an MTS server cannot exist across calls.

Socket	Because the client-side of the connection is outbound, the <code>Socket</code> instance can be serviced across calls within either an MTS or dedicated server.
--------	--

`ServerSocket` The server-side of the connection is a listener.

- **Dedicated server**—Your `ServerSocket` can listen across calls only within a dedicated server; the dedicated server exists solely for servicing the single client.
- **MTS server**—The `ServerSocket` is closed at the end of a call within an MTS server; the MTS uses shared servers, which move on to another client at the end of every call. You will receive an I/O exception stating that the socket was closed if you try to use the `ServerSocket` outside of the call it was created in.

Threading in JServer

The Aurora JVM implements a non-preemptive threading model. With this model, the JVM runs all Java threads on a single operating system thread. It schedules them in a round-robin fashion and switches between them only when they block. Blocking occurs when you, for example, invoke the `Thread.yield()` method or wait on a network socket by invoking `mySocket.read()`.

Advantages of JServer's Threading Model	Disadvantages
<ul style="list-style-type: none">■ simple to program■ efficient to implement in the Java virtual machine, because a thread switch does not require any system calls■ safer, because the JVM can detect a deadlock that would hang a preemptive JVM and can then raise a runtime exception	<ul style="list-style-type: none">■ does not exhibit any concurrency■ lack of portability■ performance considerations, because of the system calls required for locking when blocking the thread■ memory scalability, because efficient multi-threaded memory allocation requires a larger pool of memory

Oracle chose this model because any Java application written on a single-processor system works identical to one written on a multi-processor system. Also, the lack of concurrency among Java threads is not an issue because Aurora is embedded in the database, which provides a higher degree of concurrency than any conventional JVM.

There is no need to use threads within the application logic because the Oracle server preemptively schedules the session JVMs. If you must support hundreds or thousands of simultaneous transactions, start each one in its own JVM. This is exactly what happens when you create a session on the JServer. The normal

transactional capabilities of the Oracle database server accomplish coordination and data transfer between the Java virtual machines. This is not a scalability issue, because in contrast to the 6 to 8 MB memory footprint of the typical Java virtual machine, the Oracle server can create thousands of Java virtual machines, with each one taking less than 40 KB.

Threading is managed within Aurora by servicing a single thread until it completes or blocks. If the thread blocks, by yielding or waiting on a network socket, the JVM will service another thread. However, if the thread never blocks, it is serviced until completed.

The Aurora JVM has added the following features for better performance and thread management:

- System calls are at a minimum. Aurora has exchanged some of the normal system calls with non-system solutions. For example, entering a monitor-synchronized block or method does not require a system call.
- Deadlocks are detected.
 - * Aurora monitors for deadlocks between threads. If a deadlock occurs, Aurora terminates one of the threads and throws the `oracle.aurora.vm.DeadlockError` exception.
 - * Single-threaded applications cannot suspend. If the application has only a single thread and you try to suspend it, the `oracle.aurora.vm.LimboError` exception is thrown.

Thread Lifecycle

In the single-threaded execution case, the call ends when one of the following events occurs:

1. The thread returns to its caller.
2. An exception is thrown and is not caught in Java code.
3. The `System.exit()`, `oracle.aurora.vm.OracleRuntime.exitCall()`, or `oracle.aurora.mts.session.Session.THIS_SESSION().endSession()` method is invoked.

If the initial thread creates and starts other Java threads, the rules about when a call ends are slightly more complicated. In this case, the call ends in one of the following two ways:

1. The main thread returns to its caller, or an exception is thrown and not caught in this thread, *and* all other non-daemon threads complete execution.

Non-daemon threads complete either by returning from their initial method or because an exception is thrown and not caught in the thread.

2. Any thread invokes the `System.exit()`, `oracle.aurora.vm.OracleRuntime.exitCall()`, or `oracle.aurora.mts.session.Session.THIS_SESSION().endSession()` method.

When a call ends because of a return and/or uncaught exceptions, Aurora throws a `ThreadDeathException` in all daemon threads. The `ThreadDeathException` essentially forces threads to stop execution.

When a call ends because of a call to `System.exit()`, `oracle.aurora.vm.OracleRuntime.exitCall()`, or `oracle.aurora.vm.oracleRuntime.exitSession()`, Aurora ends the call abruptly and terminates all threads, but does not throw `ThreadDeathException`.

During the execution of a single call, a Java program can recursively cause more Java code to be executed. For example, your program can issue a SQL query using JDBC or SQLJ that in turn causes a trigger written in Java to be invoked. All the preceding remarks regarding call lifetime apply to the top-most call to Java code, not to the recursive call. For example, a call to `System.exit()` from within a recursive call will exit the entire top-most call to Java, not just the recursive call.

Invoking Java in the Database

We reviewed the basics of writing and deploying Java applications on Oracle8i in [Chapter 2, "Writing Java Applications on Oracle8i"](#). This chapter gives you an overview and examples for how to invoke Java within the database.

- [Overview](#)
- [Invoking Java Methods](#)
- [Utilizing SQLJ and JDBC for Querying Database](#)
- [Debugging Server Applications](#)
- [How To Tell You Are Executing in the Server](#)
- [Redirecting Output on the Server](#)

Overview

In Oracle8i, you utilize Java in one of the following ways:

- **Invoking Java Methods**—Invoke Java methods in classes loaded within the database. This includes Java stored procedures, Servlets, JavaServer Pages, CORBA, and EJB.

Note: For more information on JavaServer Pages and Servlets, see the *Oracle JavaServer Pages Developer's Guide* and *Oracle8i Oracle Servlet Engine User's Guide*.

- **Utilizing SQLJ and JDBC for Querying Database**—You can query the database from a Java client through utilizing JDBC or SQLJ.

We recommend that you approach Java development in Oracle8i incrementally, building on what you learn at each step. The easiest way to invoke Java within the database is through Java stored procedures. Once you have mastered that, you should move on to CORBA and EJB applications.

1. You should master the process of writing simple Java stored procedures as listed in "[Preparing Java Class Methods for Execution](#)" on page 2-15 and the *Oracle8i Java Stored Procedures Developer's Guide*. This includes writing the Java class, deciding on a resolver, loading the class into the database, and publishing the class.
2. You should understand how to access and manipulate SQL data from Java. Most Java server programs, and certainly Java programs executing on Oracle8i, interact with database-resident data. The two standard APIs for accomplishing this are JDBC and SQLJ. Because JDBC forms the foundation for SQLJ, you should understand how the two work together, even though you might be using only SQLJ in your code.
3. If you intend to distribute Java logic between client and server or in an N-tier architecture, you should understand how CORBA and EJB work in Oracle8i. CORBA and EJB provide the simplest solution to this difficult problem, in an Internet-standard manner, enabling you to leverage component-based development for transactional applications. Furthermore, EJB and CORBA utilize Oracle8i's facilities for Java stored procedures and JDBC.

Java is a simple, general purpose language for writing stored procedures. JDBC and SQLJ allow Java to access SQL data. They support SQL operations and concepts, variable bindings between Java and SQL types, and classes that map Java classes to

SQL types. You can write portable Java code that can execute on a client or a server without change. With JDBC and SQLJ, the dividing line between client and server is usually obvious—SQL operations happen in the server, and application program logic resides in the client.

As you write more complex Java programs, you can gain performance and scalability by controlling the location where the program logic executes. You can minimize network traffic and maximize locality of reference to SQL data. JDBC and SQLJ furnish ways to accomplish these goals. However, as you tend to leverage the object model in your Java application, a more significant portion of time is spent in Java execution, as opposed to SQL data access and manipulation. It becomes more important to understand and specify where Java objects reside and execute in an Internet application. Now you have become a candidate for moving into the world of CORBA and Enterprise JavaBeans.

Invoking Java Methods

The way your client calls a Java method depends on the type of Java application. The following sections discuss each of the Java APIs available for creating a Java class that can be loaded into the database and accessed by your client:

- [Utilizing Java Stored Procedures](#)
- [Utilizing Distributed Objects With CORBA and EJB](#)
- [Utilizing Remote Method Invocation \(RMI\)](#)
- [Utilizing Java Native Interface \(JNI\) Support](#)
- [Utilizing SQLJ and JDBC for Querying Database](#)

Utilizing Java Stored Procedures

You execute Java stored procedures similarly to PL/SQL. Normally, calling a Java stored procedure is a by-product of database manipulation in that it is usually the result of a trigger or DML call.

To invoke a Java stored procedure, you must publish it through a call specification. The following example shows how to create, resolve, load, and publish a simple Java stored procedure that echoes “Hello world”.

1. Write the Java class.

Define a class, `Hello`, with one method, `Hello.world()`, that returns the string “Hello world”.

```
public class Hello
{
    public static String world ()
    {
        return "Hello world";
    }
}
```

2. Compile the class on your client system. Using Sun Microsystem's JDK, for example, you invoke the Java compiler, `javac`, as follows:

```
javac Hello.java
```

Normally, it is a good idea to specify your `CLASSPATH` on the `javac` command line, especially when writing shell scripts or make files. The Java compiler produces a Java binary file—in this case, `Hello.class`.

Keep in mind where this Java code will execute. If you execute `Hello.class` on your client system, it searches the `CLASSPATH` for all supporting core classes it must execute. This search should result in locating the dependent class in one of the following:

- as an individual file in a directory, where the directory is specified in the `CLASSPATH`
 - within a `.jar` or `.zip` file, where the directory is specified in the `CLASSPATH`
3. Decide on the resolver for your class.

In this case, you load `Hello.class` in the server, where it is stored in the database as a Java schema object. When you execute the `world()` method of the `Hello.class` on the server, it finds the necessary supporting classes, such as `String`, using a resolver—in this case, the default resolver. The default resolver looks for classes in the current schema first and then in `PUBLIC`. All core class libraries, including the `java.lang` package, are found in `PUBLIC`. You may need to specify different resolvers, and you can force resolution to occur when you use `loadjava`, to determine if there are any problems earlier, rather than at runtime. Refer to ["Resolving Class Dependencies"](#) on page 2-19 or the *Oracle8i Java Tools Reference* for more details on resolvers and `loadjava`.

4. Load the class on the Oracle8i server using `loadjava`. You must specify the username and password.

```
loadjava -user scott/tiger Hello.class
```


5. Publish the stored procedure through a call specification

To invoke a Java static method with a SQL CALL, you must publish it with a call specification. A call specification defines for SQL which arguments the method takes and the SQL types it returns.

In SQL*Plus, connect to the database and define a top-level call specification for `Hello.world()`:

```
SQL> connect scott/tiger
connected
SQL> create or replace function HELLOWORLD return VARCHAR2 as
  2 language java name 'Hello.world () return java.lang.String';
  3 /
Function created.
```

6. Invoke the stored procedure

```
SQL> variable myString varchar2[20];
SQL> call HELLOWORLD() into :myString;
Call completed.
SQL> print myString;
```

```
MYSTRING
```

```
-----
Hello world
```

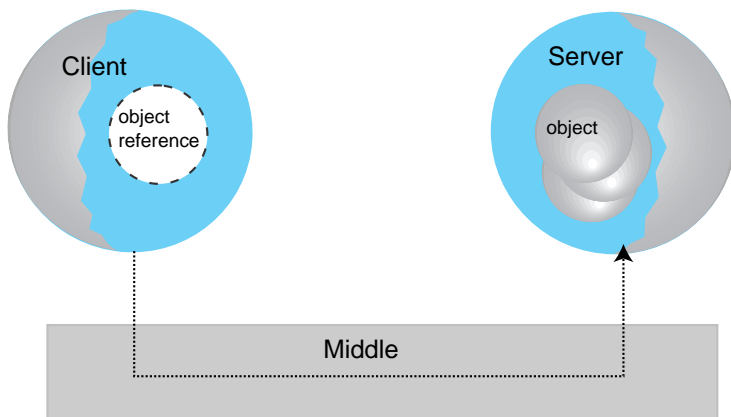
```
SQL>
```

The call `HELLOWORLD() into :myString` statement performs a top-level call in Oracle8i. The Oracle-specific `select HELLOWORLD from DUAL` also works. Note that SQL and PL/SQL see no difference between a stored procedure written in Java, PL/SQL, or any other language. The call specification provides a means to tie inter-language calls together in a consistent manner. Call specifications are necessary only for entry points invoked with triggers or SQL and PL/SQL calls. Furthermore, JDeveloper can automate the task of writing call specifications.

For more information on Java stored procedures, using Java in triggers, call specifications, rights models, and inter-language calls, refer to the *Oracle8i Java Stored Procedures Developer's Guide*.

Utilizing Distributed Objects With CORBA and EJB

In a program whose logic is distributed, the architecture of choice has three tiers—the client, the middle tier, and the database server.



- | | |
|-------------|--|
| Client tier | Typically limited to display of information provided by the middle tier. |
| Middle tier | Facilitates the communication between client and server. Typically manages the server objects. Marshals and unmarshals the parameters and return values. |
| Server tier | Performs the business or application logic. |

The server object within the three-tier model does the business logic. This may or may not include accessing a database for SQL queries. Oracle8i removes the need for a physical middle tier for distributed applications where the server object requires access to a database. Oracle8i still maintains a three-tier logical architecture, but by combining the middle tier and the database server, the physical architecture is two-tier. The flexibility inherent in this architecture is ideally suited to Internet applications where the client presents information in a Web browser, interacting with servers across the network. Those servers, in turn, can be federated and cooperate in their own client-server interactions to provide information to Web-based clients in an intranet or Internet application.

To use the two-tier distributed object approach for your application, you can use either the CORBA or EJB APIs.

- CORBA uses Interface Definition Language (IDL) to specify, in a language-independent manner, how to access and use a group of objects known

as a component. Oracle8i interacts with each client as if it had its own Java virtual machine running in the server. There is no single ORB in the JServer servicing multiple client requests. Instead, JServer leverages off of Oracle8i's Multithreaded Server (MTS) architecture, providing an ORB per session.

- Enterprise JavaBeans relies on the following:
 - * Java class definitions specify the interface to a component.
 - * RMI-style declarative deployment descriptors define how the component is treated in a transactional, secure application.

An EJB programmer writes business logic and the interfaces to the component; a deployment tool, `deployejb`, loads and publishes the component. No knowledge of IDL is necessary. This portable Java-based server framework provides a fast, scalable, and easy solution to Java-based, three-tier applications.

CORBA and EJB are complementary. The JServer implementation of the Enterprise JavaBeans 1.1 specification builds on the underlying support and services of CORBA.

IIOP Transport

Unlike a session in which the client communicates through Net8, you access CORBA and EJB sessions through IIOP, which is capable of servicing multiple client connections. Although scalable applications generally provide one session per client-server interaction, the ability to service multiple clients extends the flexibility of the session. IIOP enables callouts, callbacks, and loopbacks in your distributed communications.

Naming

You can access components through a name service, which forms a tree, similar to a file system, where you can store objects by name. When you put a CORBA or EJB object into the namespace, you are publishing it.

There are two supported naming protocols within Oracle8i:

- Java Naming and Directory Interface (JNDI)—This package provides a unified interface to name services. Part of JNDI provides a platform-independent abstraction for accessing a file system, which is platform dependent.
- CORBA's CosNaming—If you are a CORBA programmer, you are familiar with bootstrapping your application using CosNaming. Your client Java code obtains handles to objects that reside on the server. Those objects are reachable through the name service. The ORB supplies the name service, which presupposes that

the ORB is running when you attempt to locate the server objects when bootstrapping your client application. JServer provides an activation service based on CORBA's CosNaming. You use a URL-based name within JNDI as an interface to CosNaming when referring and activating CORBA objects in a session. This namespace incorporates the idea of a session directly in the URL, allowing the client to easily manipulate multiple sessions. All bootstrapping is performed by establishing a session with the JServer and using objects always reachable from the Oracle8i database that the standard JNDI and CORBA CosNaming make visible to you. You do not use Inter-ORB References (IORs) as with most CORBA applications.

Creating and Deploying Enterprise JavaBeans

CORBA and EJB application development are complicated topics covered in the *Oracle8i Enterprise JavaBeans Developer's Guide and Reference* and the *Oracle8i CORBA Developer's Guide and Reference*. This section gives an example of how to create an EJB component. This example creates an EmployeeBean to look up an employee record in the Oracle RDBMS.

1. Create the home interface. The home interface will reside in the server, enabling you to create instances of your EJB on the server. A home interface is a Java interface that extends EJBHome. The home interface is the only object published in the namespace to be visible to clients. The client can access the component through JNDI.
2. Create the remote interface. The remote interface specifies the methods you implement in the EJB, such as instance methods you can invoke from a client. A remote interface is a Java interface that extends EJBObject. As an interface, you use it to specify the methods implemented in the bean.
3. Implement the bean class and the methods that the remote interface defines. In the EmployeeBean example, the only method is `getEmployee()`. You write a bean by creating a class that implements the SessionBean interface.
4. Create the deployment descriptor. The deployment descriptor specifies attributes of the bean, including its transactional properties and security treatment. You specify the attributes, and the `deployejb` tool ensures that the server enforces them.
5. Deploy the EJB. When you deploy the EJB, the `deployejb` tool does the following:
 - Places your home interface and the EJB on the server.
 - Publishes the home interface in the namespace.

- Generates the Java code on the server side to manage transactions and security specified in the deployment descriptor.
- Generates and returns the stub interfaces that provide the client access to the remote functionality of the bean.

Using an EJB

Once you create and deploy an EJB, you will want to use it from a client program. You can use EJBs between servers in n-tier applications also, in which case the client for one server can also be the server for other clients. In your client code, you must perform the following steps:

1. Locate the home interface object that resides on the server. You will locate the object using Java-standard JNDI lookup facilities.
2. Authenticate the client to the server. EJB and CORBA clients use database sessions, just as with any other Oracle client. To initiate a session, you must let the server know you are a valid user. You can use several different approaches to accomplish authentication in a secure manner.
3. Activate an instance of the bean. Because the object you locate with JNDI is the home interface, you will use one of its `create()` methods to return an activated instance of the EJB.
4. Invoke methods on the bean. When you invoke a method on the bean, the method is actually executed in the server, and the appropriate parameter and return objects are transparently transported (by copy) across the underlying IIOP connection. All objects the EJBs return must be serializable—they must implement `java.io.Serializable`.

The *Oracle8i Enterprise JavaBeans Developer's Guide and Reference* and the *Oracle8i CORBA Developer's Guide and Reference* books discuss the details of these steps. Java IDEs, as with Oracle's JDeveloper, can automate and simplify the deployment and descriptor process.

Session Shell

Session shell is an example of a tool written completely in Java using Java stored procedures and CORBA. It interacts with server-resident objects that are visible through CORBA within your session by using UNIX shell commands. For more information, see the *The Oracle8i Enterprise JavaBeans Developer's Guide and Reference* and the *Oracle8i CORBA Developer's Guide and Reference* books. This tool demonstrates how you can use CORBA to build tools that make life simpler for developers and end users.

The session shell provides a shell-like interface to the server. This shell allows users to manipulate the session namespace with familiar UNIX commands, such as `mkdir`, `ls`, and `rm`. In addition, the session shell furnishes a convenient way to run Java programs in the server, using the `java` command. The session shell `java` command takes the name of a class and any arguments the user types in. The session shell calls the static `main(String[])` method on the class, running the Java program in the server. `System.out` and `System.err` are captured and transparently redirected back to the user's console.

Utilizing Remote Method Invocation (RMI)

JServer fully supports Java Remote Method Invocation (RMI). All RMI classes and `java.net` support are in place. In general, RMI is not useful or scalable in JServer applications. CORBA and EJB are the preferred APIs for invoking methods of remote objects. The RMI Server that Sun Microsystems supplies does function on the JServer platform. Because Sun Microsystems's RMI Server uses operating system sockets and is not accessible through a presentation, it is useful only within the context of a single call. It relies heavily on Java language level threads. By contrast, the Oracle8i ORB and EJB rely on the database server to gain scalability. You can efficiently implement an RMI server as a presentation; however, CORBA and EJB currently serves this purpose.

Note: A presentation is an object that accepts either a Net8 or IIOP incoming connection into the database. See "[Configuring JServer](#)" on page 4-6 for more information.

Utilizing Java Native Interface (JNI) Support

The Java Native Interface (JNI) is a standard programming interface for writing Java native methods and embedding the Java virtual machine into native applications. The primary goal of JNI is to provide binary compatibility of Java applications that use platform-specific native libraries.

Oracle does not support the use of JNI in JServer applications. If you use JNI, your application is not 100% pure Java, and the native methods require porting between platforms. Native methods have the potential for crashing the server, violating security, and corrupting data.

Utilizing SQLJ and JDBC for Querying Database

You can use one of two protocols for querying the database from a Java client. Both protocols establish a session with a given username/password to the database and execute SQL queries against the database.

JDBC Use this protocol for more complex or dynamic SQL queries. JDBC requires you to establish the session, construct the query, and so on.

SQLJ Use this protocol for static, easy SQL queries. SQLJ is typically a one-liner that executes against a known table with known column names.

JDBC

JDBC is an industry-standard API developed by Sun Microsystems that allows you to embed SQL statements as Java method arguments. JDBC is based on the X/Open SQL Call Level Interface and complies with the SQL92 Entry Level standard. Each vendor, such as Oracle, creates its JDBC implementation by implementing the interfaces of the Sun Microsystems `java.sql` package. Oracle offers three JDBC drivers that implement these standard interfaces:

1. The JDBC Thin driver, a 100% pure Java solution you can use for either client-side applications or applets and requires no Oracle client installation.
2. The JDBC OCI drivers (OCI 8 or OCI 7), which you use for client-side applications and requires an Oracle client installation.
3. The server-side JDBC driver embedded in the Oracle8i server.

For the developer, using JDBC is a step-by-step process of creating a statement object of some type for your desired SQL operation, assigning any local variables that you want to bind to the SQL operation, and then executing the operation. This process is sufficient for many applications but becomes cumbersome for any complicated statements. Dynamic SQL operations, where the operations are not known until runtime, require JDBC. In typical applications, however, this represents a minority of the SQL operations.

SQLJ

SQLJ offers an industry-standard way to embed any static SQL operation directly into Java source code in one simple step, without requiring the individual steps of JDBC. Oracle SQLJ complies with ANSI standard X3H2-98-320.

SQLJ consists of a translator—a precompiler that supports standard SQLJ programming syntax—and a runtime component. After creating your SQLJ source code in a `.sqlj` file, you process it with the translator, which translates your SQLJ source code to standard Java source code, with SQL operations converted to calls to the SQLJ runtime. In the Oracle SQLJ implementation, the translator invokes a Java compiler to compile the Java source. When your Oracle SQLJ application runs, the SQLJ runtime calls JDBC to communicate with the database.

SQLJ also allows you to catch errors in your SQL statements before runtime. JDBC code, being pure Java, is compiled directly. The compiler has no knowledge of SQL, so it is unaware of any SQL errors. By contrast, when you translate SQLJ code, the translator analyzes the embedded SQL statements semantically and syntactically, catching SQL errors during development, instead of allowing an end-user to catch them when running the application.

An Example Comparing JDBC and SQLJ

The following is an example of a simple operation, first in JDBC code and then SQLJ code.

JDBC:

```
// (Presume you already have a JDBC Connection object conn)
// Define Java variables
String name;
int id=37115;
float salary=20000;

// Set up JDBC prepared statement.
PreparedStatement pstmt = conn.prepareStatement
    ("select ename from emp where empno=? and sal>?");
pstmt.setInt(1, id);
pstmt.setFloat(2, salary);

// Execute query; retrieve name and assign it to Java variable.
ResultSet rs = pstmt.executeQuery();
while (rs.next()) {
    name=rs.getString(1);
    System.out.println("Name is: " + name);
}

// Close result set and statement objects.
rs.close();
pstmt.close();
```


1. Define the Java variables `name`, `id`, and `salary`.
2. Define a prepared statement (this presumes you have already established a connection to the database so that you can use the `prepareStatement()` method of the connection object).

You can use a prepared statement whenever values within the SQL statement must be dynamically set (you can use the same prepared statement repeatedly with different variable values). The question marks in the prepared statement are placeholders for Java variables and are given values in the

`pstmt.setInt()` and `pstmt.setFloat()` lines of code. The first “?” is set to the `int` variable `id` (with a value of 37115). The second “?” is set to the `float` variable `salary` (with a value of 20000).

3. Execute the query and return the data into a JDBC result set object. (You can use result sets to gather query data.)
4. Retrieve the data of interest (the name) from the result set and print it. A result set usually contains multiple rows of data, although this example has only one row.

By comparison, here is some SQLJ code that performs the same task. Note that all SQLJ statements, both declarations and executable statements, start with the `#sql` token.

SQLJ:

```
String name;
int id=37115;
float salary=20000;
#sql {select ename into :name from emp where empno=:id and sal>:salary};
System.out.println("Name is: " + name);
```

SQLJ, in addition to allowing SQL statements to be directly embedded in Java code, supports Java host expressions (also known as bind expressions) to be used directly in the SQL statements. In the simplest case, a host expression is a simple variable as in this example, but more complex expressions are allowed as well. Each host expression is preceded by “:” (colon). This example uses Java host expressions `name`, `id`, and `salary`. In SQLJ, because of its host expression support, you do not need a result set or equivalent when you are returning only a single row of data.

Complete SQLJ Example

This section presents a complete example of a simple SQLJ program:

```
import java.sql.*;
import sqlj.runtime.ref.DefaultContext;
import oracle.sqlj.runtime.Oracle;
#sql iterator MyIter (String ename, int empno, float sal);

public class MyExample
{
    public static void main (String args[]) throws SQLException
    {
        Oracle.connect
            ("jdbc:oracle:thin:@oow11:5521:sol2", "scott", "tiger");

        #sql { insert into emp (ename, empno, sal)
            values ('SALMAN', 32, 20000) };
        MyIter iter;

        #sql iter={ select ename, empno, sal from emp };
        while (iter.next()) {
            System.out.println
                (iter.ename()+" "+iter.empno()+" "+iter.sal());
        }
    }
}
```

1. Declare your iterators. SQLJ uses a strongly typed version of JDBC result sets, known as iterators. The main difference between the two is that an iterator has a specific number of columns of specific datatypes. You must define your iterator types beforehand, as in this example:

```
#sql iterator MyIter (String ename, int empno, float sal);
```

This declaration results in SQLJ creating an iterator class `MyIter`. Iterators of type `MyIter` can store results whose first column maps to a Java `String`, whose second column maps to a Java `int`, and whose third column maps to a Java `float`. This definition also names the three columns—`ename`, `empno`, and `sal`, respectively—to match the table column names in the database. `MyIter` is a named iterator. See Chapter 3 of the *Oracle8i SQLJ Developer's Guide and Reference* to learn about positional iterators, which do not require column names.

2. Connect to the database.

```
Oracle.connect("jdbc:oracle:thin:@oow11:5521:sol2","scott", "tiger");
```

Oracle SQLJ furnishes the `Oracle` class, and its `connect()` method accomplishes three important things:

- a. Registers the Oracle JDBC drivers SQLJ uses to access the database.
- b. Opens a database connection for the specified schema (user `scott`, password `tiger`) at the specified URL (host `oow11`, port `5521`, SID `sol2`, “thin” JDBC driver).
- c. Establishes this connection as the default connection for your SQLJ statements. Although each JDBC statement must explicitly specify a connection object, a SQLJ statement can either implicitly use a default connection or optionally specify a different connection.

3. Execute a SQL statement.

- a. Insert a row into the `emp` table:

```
#sql {insert into emp (ename, empno, sal) values ('SALMAN', 32, 20000)};
```

- b. Instantiate and populate the iterator:

```
MyIter iter;
#sql iter={select ename, empno, sal from emp};
```

4. Access the data that was populated within the iterator.

```
while (iter.next()){
    System.out.println(iter.ename()+" "+iter.empno()+" "+iter.sal());
}
```

The `next()` method is common to all iterators and plays the same role as the `next()` method of a JDBC result set, returning `true` and moving to the next row of data if any rows remain. You access the data in each row by calling iterator accessor methods whose names match the column names (this is a characteristic of all named iterators). In this example, you access the data using the methods `ename()`, `empno()`, and `sal()`.

SQLJ Strong Typing Paradigm

SQLJ uses strong typing—such as iterators—instead of result sets, which allows your SQL instructions to be checked against the database during translation. For example, SQLJ can connect to a database and check your iterators against the database tables that will be queried. The translator will verify that they match, allowing you to catch SQL errors during translation that would otherwise not be caught until a user runs your application. Furthermore, if changes are subsequently

made to the schema, you can determine if this affects the application simply by re-running the translator.

Translating a SQLJ Program

Integrated development environments such as Oracle JDeveloper, a Windows-based visual development environment for Java programming, can translate, compile, and customize your SQLJ program for you as you build it. If you are not using an IDE, then you can use the front-end SQLJ utility, `sqlj`. You can run it as follows:

```
%sqlj MyExample.sqlj
```

The SQLJ translator checks the syntax and semantics of your SQL operations. You can enable online checking to check your operations against the database. If you choose to do this, you must specify an example database schema in your translator option settings. It is not necessary for the schema to have identical data to the one the program will eventually run against; however, the tables should have columns with corresponding names and datatypes. Use the `user` option to enable online checking and specify the username, password, and URL of your schema, as in the following example:

```
%sqlj -user=scott/tiger@jdbc:oracle:thin:@oow11:5521:sol2 MyExample.sqlj
```

Running a SQLJ Program in the Server

Many SQLJ applications run on a client; however, SQLJ offers an advantage in programming stored procedures—which are usually SQL-intensive—to run in the server.

There is almost no difference between coding for a client-side SQLJ program and a server-side SQLJ program. The SQLJ runtime packages are automatically available on the server, and there are just the following few considerations:

- There are no explicit database connections for code running in the server; only a single implicit connection. You do not need the usual connection code. If you are porting an existing client-side application, you do not have to remove your connection code, because it will be ignored.
- The JDBC server-side internal driver does not support auto-commit functionality. Use SQLJ syntax for manual commits and rollbacks of your transactions.
- On the server, the default output device is a trace file, not the user screen. This is normally an issue or question only for development because you would not write to `System.out` in a deployed server application.

To run a SQLJ program in the server, presuming you developed the code on a client, you have two options:

- Translate your SQLJ source code on the client and load the individual components (Java classes and resources) to the server. In this case, it is easiest to bundle them into a `.jar` file first.
- Load your SQLJ source code to the server for the embedded translator to translate.

In either case, you can use the Oracle `loadjava` utility to load the file or files to the server. See the *Oracle8i SQLJ Developer's Guide and Reference* for more information.

Converting a Client Application to Run in the Server

The steps in converting an existing SQLJ client-side application to run in the server are as follows. Assume this is an application that has already been translated on the client:

1. Create a `.jar` file for your application components.
2. Use the `loadjava` utility to load the `.jar` file to the server.
3. Create a SQL wrapper in the server for your application. For example, to run the preceding `MyExample` application in the server:

```
create or replace procedure SQLJ_MYEXAMPLE as language java
  name 'MyExample.main(java.lang.String[])';
```

You can then execute `SQLJ_MYEXAMPLE` as with any other stored procedure.

Interacting with PL/SQL

All the Oracle JDBC drivers communicate seamlessly with Oracle SQL and PL/SQL, and it is important to note that SQLJ interoperates with PL/SQL. You can start using SQLJ without having to rewrite any PL/SQL stored procedures. Oracle SQLJ includes syntax for calling PL/SQL stored procedures and also allows PL/SQL anonymous blocks to be embedded in SQLJ executable statements, just as with SQL operations.

Debugging Server Applications

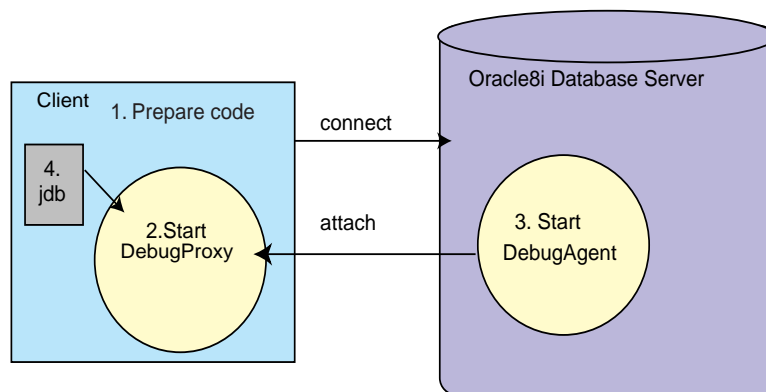
JServer furnishes a debugging capability useful for developers who use the JDK's `jdb` debugger. Oracle's JDeveloper provides a user-friendly integration with this JServer debugging feature. See the JDeveloper documentation for more information on how to debug your Java application through JDeveloper. Other independent IDE vendors will be able to integrate their own debuggers with JServer.

The Sun Microsystems `jdb` debugger attaches itself to an executing process, and helps you debug the executing process. The application that you are debugging must have been compiled with the debug option (`-g`).

In JServer, your Java program executes remotely on a server. The server can reside on the same physical machine, but it typically resides on a separate machine. Oracle8i provides a method for `jdb` to debug a Java application loaded into JServer.

This method involves an debug agent that is executing on the Oracle8i server and communicating with the executing Java application, a debug proxy that exists on the client and communicates with the Oracle8i server, and a way for `jdb` to attach itself to the debug proxy. [Figure 3-1](#) shows the relationship between the debug agent, the debug proxy, and the `jdb` debugger.

Figure 3-1 *Debug Proxy and Debug Class Facilitate jdb Debugger*



As shown in [Figure 3-1](#), the steps for remotely debugging your Java application are as follows:

1. Prepare your code for debugging.

2. Start the `DebugProxy`. The `DebugProxy` waits for a `DebugAgent` to attach to it from the server.
3. Start the `DebugAgent` giving it the debug proxy address. This starts the communication between the debug agent and the debug proxy.
4. Attach the `jdb` debugger to the debug proxy. Once attached, you can use the regular `jdb` commands.

1. Prepare the Code for Debugging

The code must be compiled with the `-g` option and the source must be made available for the debug agent to locate.

You can cause your application to be compiled with the debug option (`-g`) in one of the two following ways:

- Inform the server to compile the class with the debug option through the `set_compiler_option` procedure, as follows:

```
SQL> call dbms_java.set_compiler_option('myPackage.myCode','debug','true');
```

Note: The `set_compiler_option` procedure specifies many different compiler options on a certain class, package, or all classes. This example shows setting the option for a single class. See the `loadjava` section of the *Oracle8i Java Tools Reference* for more information on this procedure.

Then, you must load the source code using `loadjava`, as follows:

```
% loadjava -u SCOTT/TIGER -v -f -r myCode.java
```

The server will compile this class with the debug option. Also, the server now has access to both the source and the compiled binary, which the debug agent needs for showing the breakpoints.

- Compile your code on the client with the `-g` option, load the compiled class into the server, and copy the Java source file to the file system where `Oracle8i JServer` exists, as follows:

```
% javac -g MyCode.java
% loadjava -u SCOTT/TIGER -v -f -r myCode.class
% ftp dbhost
> cd /private/sourcecode
> put myCode.java
```

When `jdb` starts, set the location of the source code with `jdb`'s `'use'` command. This enables the debug agent to find the source code.

```
> use /private/sourcecode
```

Note: In order to copy any files to the database file system, you must have the correct `FilePermission`. See the Security chapter for more information.

2. Start the Debug Proxy

The `DebugProxy` class enables your remote Java application appear to be local. The debug proxy forwards all `jdb` requests to the debug agent on the server and returns the results to the attached `jdb` debugger.

Once started, the debug proxy waits for the debug agent to attach itself. Assuming the `aurora_client.jar` file is part of your `CLASSPATH`, you start the debug proxy as follows:

```
debugproxy
```

You can also specify a particular port to wait on.

```
debugproxy -port 2286
```

The proxy prints out its name, its address, and the port it is waiting on.

```
Proxy Name: yourmachinename  
Proxy Address: aaa.bbb.ccc.ddd  
Proxy Port: 2286
```

However, the easiest method to start the `DebugProxy` is to append a command to start up the `jdb` debugger at the end of the `debugproxy` command. The `debugproxy` command takes in any option given, beyond the optional port, as a command to execute after it has started. If you choose this method, you do not need to execute step 4.

For UNIX, provide the following within an executable shell script called `startjdb`:

```
#!/bin/sh  
xterm -e jdb -password &1 &
```

Then, you can automatically start up the `jdb` debugger within the `debugproxy` command, as follows:


```
debugproxy -port 1638 startjdb
```

For all Windows NT environments, provide the following within a batch file called `startjdb.bat`:

```
start jdb -password %1
```

Then, you can automatically start up the `jdb` debugger within the `debugproxy` command, as follows:

```
debugproxy -port 1638 startjdb.bat
```

3. Starting, Stopping, and Restarting the Debug Agent

After you connect to the server (starting a session) and start a debug proxy, you can start a debug agent on the server that will connect to the proxy. When the `DebugAgent` starts, the `DebugProxy` displays a password to use when attaching the debugger in step 4.

Note: You must have the debug permission, `JAVADEBUGPRIV`, granted to your user to run a debug agent. See "[Debugging Permissions](#)" in the *Oracle8i Java Developer's Guide* for more information.

Once a proxy is running, you can start a debug agent to connect to the proxy from `SQL*Plus`. You must specify the IP address or URL for a machine running a debug proxy, the port the proxy is waiting on, and a timeout in seconds. You start and stop the debug agent using methods specified within the `DBMS_JAVA` package.

```
SQL> call dbms_java.start_debugging('yourmachinename', 2286, 66);
```

There is no way to cause server-resident code to execute and break, that is, execute and remain indefinitely in a halted mode. Instead, when you start the `DebugAgent`, you must specify a timeout period for the `DebugAgent` to wait before terminating. The start call waits until the timeout expires or until the main thread is suspended and resumed before it completes. Calculate a timeout that includes enough time for your debugger to start up, but not so much as to delay your session if you cannot connect a debugger.

Note: If an agent is already running, Aurora stops it and starts a new agent.

You can stop the debug agent explicitly through the `stop_debugging` method.

```
SQL> call dbms_java.stop_debugging();
```

Once a debug agent starts, it runs until you stop it, the debugger disconnects, or the session ends.

You can restart a stopped agent with any breakpoints still set with the `restart_debugging` method. The call waits until the timeout expires before it completes. You can also restart a running agent just to buy some seconds to suspend threads and set breakpoints.

```
SQL> call dbms_java.restart_debugging(66);
```

OracleAgent Class

The `DBMS_JAVA` debug agent and proxy calls are published entry points to static methods that reside in `oracle.aurora.debug.OracleAgent` class. You can start, stop, and restart the debug agent in Java code using the class `oracle.aurora.debug.OracleAgent` directly through the following methods:

```
public static void start(String host, int port, long timeout_seconds);  
public static void stop();  
public static void restart(long timeout_seconds);
```

4. Connecting a Debugger

Start `jdb` and attach it to the debug proxy using the password provided by the `DebugProxy` when the `DebugAgent` connected to it. In order to preserve your timeout, suspend all threads through `jdb`, set your breakpoints, and then resume.

Each time a debug agent connects to a debug proxy, the debug proxy starts a thread to wait for connections from a debugger. The thread prints out the number, name and address of the connecting agent, the port it is waiting on, and the port encoded as a password. Here, a specific port and password are provided for illustration only:

```
Agent Number: 1  
Agent Name: servername  
Agent Address: eee.fff.jjj.kkk  
Agent Port: 2286  
Agent Password: 3i65bn
```

You can then pass the password to a `jdb`-compatible debugger (JDK 1.1.6 or later):

```
jdb -password 3i65bn
```

The first thing you should do in the debugger is suspend all threads. Otherwise, your `start_debugging` call might time out and complete before you get your breakpoints set.

If your code writes to `System.out` or `System.err`, then you may also want to use the `dbgtrace` flag to `jdb`, which redirects these streams to the debugging console:

```
jdb -dbgtrace -password 3i65bn
```

Example 3-1 Starting a DebugAgent on the Server

The following example shows how to debug an object that exists on the server. First, you need to start a proxy through the `debugproxy` command-line tool. This example starts up the proxy on the server, `tstHost`, and informs the `debugproxy` to start up the `jdb` debugger when contacted by the debug agent.

In another window, make sure that the debug agent user has the correct privileges and then start up the debug agent. Once the agent starts, the `debugproxy` starts up the `jdb` debugger and allows you to set your breakpoints. Since you have a specified amount of time before the agent times out, the first thing you should do is suspend all threads. Then, set all of your breakpoints before resuming. This suspends the timeout until you are ready to execute.

window 1 on `tstHost`

```
SQL> call dbms_java.set_compiler_option('', 'debug', 'true');
SQL> exit
% loadjava -u SCOTT/TIGER -v -f -r myCode.java
% debugproxy -port 2286 start jdb -password
. (wait until a debug agent starts up and
. contact this proxy... when it does, jdb
. starts up automatically and you can set
. breakpoints and debug the object, as follows:)
> suspend
> load SCOTT:myCode
> stop in myCode:updateAccount
> resume
> ...
```

window 2 on `tstHost`

```
SQL> grant JavaDebugPriv to SCOTT
SQL> call dbms_java.start_debugging('tstHost', 2286, 30);
```

How To Tell You Are Executing in the Server

You might want to write Java code that executes in a certain way in the server and another way on the client. In general, Oracle does not recommend this. In fact, JDBC and SQLJ go to some trouble to enable you to write portable code that avoids this problem, even though the drivers used in the server and client are different.

If you must determine whether your code is executing in the server, use the `System.getProperty` method, as follows:

```
System.getProperty ("oracle.jserver.version")
```

The `getProperty` method returns the following:

- If executing in the server, returns a `String` that represents the Oracle8i database version ("8.1.5" or "8.1.6").
- If executing on the client, returns `null`.

Redirecting Output on the Server

`System.out` and `System.err` print to the current trace files. To redirect output to the SQL*Plus text buffer, use this workaround:

```
SQL> SET SERVEROUTPUT ON  
SQL> CALL dbms_java.set_output(2000);
```

The minimum (and default) buffer size is 2,000 bytes; the maximum size is 1,000,000 bytes. In the following example, the buffer size is increased to 5,000 bytes:

```
SQL> SET SERVEROUTPUT ON SIZE 5000  
SQL> CALL dbms_java.set_output(5000);
```

Output prints at the end of the call.

For more information about SQL*Plus, see the *SQL*Plus User's Guide and Reference*.

Java Installation and Configuration

This chapter details what you need to know to install and configure JServer within your Oracle8i database. To configure Java memory, see the "[Java Memory Usage](#)" section in [Chapter 6, "Oracle8i Java Application Performance"](#).

- [Initializing a Java-Enabled Database](#)
- [Configuring JServer](#)
- [Enabling the Java Client](#)

Initializing a Java-Enabled Database

If you installed Oracle8i with JServer, the database is Java-enabled. That is, it is ready to run Java stored procedures, JDBC, SQLJ, and CORBA/EJB objects. If you are using your own scripts to create your Oracle instance, you must initialize the JServer explicitly.

You install JServer in one of three ways:

- Oracle8i Typical or Minimal Install—Choose the Typical or Minimal Oracle8i installation, which results in JServer being automatically installed.
- Oracle8i Custom Install—Choose the JServer option within a "Custom" Oracle8i installation.
- **Manual Install**—Install JServer by invoking the `initjvm.sql` script.

Manual Install

If you did not install JServer through any of the Oracle8i install options, you can add JServer to an existing database with the `initjvm.sql` script in `ORACLE_HOME/javavm/install`.

The `initjvm.sql` script loads the initial set of Java classes necessary to support Java, initializes the tables for supporting Java and for the CORBA namespace, and publishes top-level entry points through call-specifications. The `initjvm.sql` script loads the support Java classes into the database, which include the following:

- the standard Java runtime
- bytecode verifier and optimizer
- Java and SQLJ compilers
- JDBC runtime
- CORBA ORB and EJB runtime
- some additional support classes, such as `DBMS_JAVA`, which are described in "[Package DBMS_JAVA](#)" on page 4-3

The `initjvm.sql` script performs the following actions:

1. Loads the classes to the SYS schema.
2. Creates public synonyms for the loaded classes to be accessible to all users.
3. Alters some of these classes to run with definer's rights to support CORBA callouts.

4. Defines database start up and shut down triggers.

Note: The `initjvm.sql` script can take up to an hour to execute.

Requirements

Initializing a Java-enabled database requires a `SHARED_POOL_SIZE` of 50 MB, a `JAVA_POOL_SIZE` of about 20 MB, an additional 30 MB of system tablespace, and enough rollback segments. If the script fails for some reason, such as a lack of resources, you can adjust resources as necessary and re-execute `initjvm.sql`. Refer to "[Java Memory Usage](#)" on page 6-7 and the `/javavm/README.txt` file for the most up-to-date information on database initialization file configuration parameters and requirements.

In addition, there are specific requirements for enabling EJB and CORBA communications. The initial settings that the Oracle8i JServer installations furnish should be sufficient to get you started. Consult the specifics of the documentation in the *Oracle8i Enterprise JavaBeans Developer's Guide and Reference*, the *Oracle8i CORBA Developer's Guide and Reference*, and the *Net8 Administrator's Guide* for more details.

Package DBMS_JAVA

When initializing the JServer, the `initjvm.sql` script creates the PL/SQL package `DBMS_JAVA`. Some entrypoints of `DBMS_JAVA` are for your use; others are only for internal use. The corresponding Java class `DbmsJava` provides methods for accessing RDBMS functionality from Java.

The `DBMS_JAVA` package supplies the following entrypoints:

```
FUNCTION longname (shortname VARCHAR2) RETURN VARCHAR2
```

Return the full name from a Java schema object. Because Java classes and methods can have names exceeding the maximum SQL identifier length, Aurora uses abbreviated names internally for SQL access. This function simply returns the original Java name for any (potentially) truncated name. An example of this function is to print the fully qualified name of classes that are invalid:

```
select dbms_java.longname (object_name) from user_objects
       where object_type = 'JAVA CLASS' and status = 'INVALID';
```

```
FUNCTION shortname (longname VARCHAR2) RETURN VARCHAR2
```

You can specify a full name to the database by using the `shortname()` routine of the `DBMS_JAVA` package, which takes a full name as input and returns the corresponding short name. This is useful when verifying that your classes loaded by querying the `USER_OBJECTS` view.

Refer to "[Shortened Class Names](#)" on page 2-32 and *Oracle8i Java Stored Procedures Developer's Guide* for examples of these functions.

```
FUNCTION get_compiler_option(what VARCHAR2, optionName VARCHAR2)
PROCEDURE set_compiler_option(what VARCHAR2, optionName VARCHAR2,
                             value VARCHAR2)
PROCEDURE reset_compiler_option(what VARCHAR2, optionName VARCHAR2)
```

These three entry points control the options of the JServer Java and SQLJ compiler Oracle8i delivers. See "[Compiling Java Classes](#)" on page 2-15 for an example of these options. Additionally, both the *Oracle8i Java Stored Procedures Developer's Guide* and the *Oracle8i SQLJ Developer's Guide and Reference* document the options and these entry points.

```
PROCEDURE set_output (buffersize NUMBER)
```

This procedure redirects the output of Java stored procedures and triggers to the `DBMS_OUTPUT` package. See "[Redirecting Output on the Server](#)" on page 3-24 for an example.

```
PROCEDURE loadjava(options varchar2)
PROCEDURE loadjava(options varchar2, resolver varchar2)
PROCEDURE dropjava(options varchar2)
```

These procedures allow you to load and drop classes within the database using a call rather than through the `loadjava` or `dropjava` command-line tools. To execute within your Java application, do the following:

```
call dbms_java.loadjava('... options...');
call dbms_java.dropjava('... options...');
```

The options are identical to those specified for the `loadjava` and `dropjava` command-line tools. Each option should be separated by a blank. You should not separate the options with a comma. The only exception for this is the `loadjava -resolver` option, which contains blanks. For `-resolver`, you should specify all other options first, separate these options by a comma, and then specify the `-resolver` option with its definition. You should not specify the following options, because they relate to the database connection for the `loadjava` command-line tool: `-thin`, `-oci8`, `-user`, `-password`. The output is directed to `stderr`.

For more information on the available options, see the *Oracle8i Java Tools Reference* for complete information on `loadjava`.

```
PROCEDURE grant_permission( grantee varchar2,
                           permission_type varchar2,
                           permission_name varchar2,
                           permission_action varchar2 )

PROCEDURE restrict_permission( grantee varchar2,
                               permission_type varchar2,
                               permission_name varchar2,
                               permission_action varchar2)

PROCEDURE grant_policy_permission( grantee varchar2,
                                   permission_schema varchar2,
                                   permission_type varchar2,
                                   permission_name varchar2)

PROCEDURE revoke_permission(permission_schema varchar2,
                            permission_type varchar2,
                            permission_name varchar2,
                            permission_action varchar2)

PROCEDURE disable_permission(key number)

PROCEDURE enable_permission(key number)

PROCEDURE delete_permission(key number)
```

These entry points control the JVM permissions. See "[Setting Permissions](#)" on page 5-6 for a description and example of these options.

```
PROCEDURE start_debugging(host varchar2, port number,
                          timeout number)

PROCEDURE stop_debugging

PROCEDURE restart_debugging(timeout number)
```

These entry points start and stop the debug agent when debugging. See "[Debugging Server Applications](#)" on page 3-18 for a description and example of these options.

Configuring JServer

When you install JServer as part of your normal Oracle8i installation, you will encounter configuration requirements for JServer within the Oracle8i Database Configuration Assistant and the Net8 Assistant. However, if you install using `initjvm.sql`, you must configure either by bringing up certain configuration tools or manually editing the initialization files.

The main configuration for Java classes within Oracle8i includes configuring Java memory requirements, the type of database processes, and the underlying connection protocol to the server.

- **Java memory requirements**—You must have at least 20 MB of `JAVA_POOL_SIZE` and 50 MB of `SHARED_POOL_SIZE`. If you installed JServer through the Oracle installer, these parameters are configured correctly. If you installed using the SQL script `initjvm.sql`, see "[Java Memory Usage](#)" on page 6-7 for information on configuring these parameters within the database initialization files.
- **Database processes**—You must decide whether to use dedicated server processes or MTS processes for your database server.
- **Connection protocol**—The presentation layer within the database defines the type of connection your client is using to access the database. In most networking protocols, the presentation layer is responsible for making sure data is represented in a format the application and session layers can accommodate. Within the database, a presentation is a service protocol that accepts incoming network requests and activates routines in the database kernel layer or in the Aurora JVM to process the requests. Currently, the two basic presentation types are Net8 and GIOP. The GIOP presentation is used for IIOP connections for EJB and CORBA applications. Most database connections default to the Net8 connection type.

Presentation protocol	Description
GIOP	Accepts IIOP or IOP over SSL requests for CORBA or EJB applications. See the <i>Oracle8i Enterprise JavaBeans Developer's Guide and Reference</i> or the <i>Oracle8i CORBA Developer's Guide and Reference</i> for configuration information.

Presentation protocol	Description
Net8	Processes incoming Net8 requests for database SQL services from Oracle tools (such as SQL*Plus) and customer-written applications (using Forms, Pro*C, or the OCI). See "Configuring Multi-Threaded Server" in Chapter 9 of the <i>Net8 Administrator's Guide</i> for configuration information.

You will require a different configuration for your database type and connection configuration, depending on the type of Java application, as listed below:

Java API	Database type	Connection configuration
Java stored procedures	Java stored procedures can run either in dedicated server mode or multi-threaded server (MTS) mode. If you are primarily developing Java stored procedures, you can run them in the dedicated server configuration.	Java clients or PL/SQL clients that trigger a Java stored procedure connect over a Net8 connection. See the <i>Net8 Administrator's Guide</i> for information on configuring a Net8 connection.
Enterprise Java Beans (EJB) or CORBA	EJB and CORBA applications run only in the MTS configuration. See <i>Net8 Administrator's Guide</i> for information on configuring MTS.	EJB and CORBA clients use the CORBA Internet Inter-Orb Protocol (IIOP). See the <i>Oracle8i Enterprise JavaBeans Developer's Guide and Reference</i> or the <i>Oracle8i CORBA Developer's Guide and Reference</i> books for information on IIOP configuration.
Both Java stored procedures and EJB or CORBA applications	If you are combining both EJB and CORBA applications with Java stored procedures in a single application, you can configure both application types as follows: <ul style="list-style-type: none"> ▪ Configure your database to support EJB and CORBA applications in an MTS configuration ▪ Configure your database to support stored procedures in a dedicated server configuration. 	You must configure both a Net8 and an IIOP connection.

Java Stored Procedure Configuration

To configure the database to run Java stored procedures, you must decide whether you want the database to run in dedicated server mode or MTS mode.

- Dedicated server mode—You must configure the database and clients, as described in the *Oracle8i Java Stored Procedures Developer's Guide*.
- MTS mode—You must configure the server for MTS mode with the MTS_DISPATCHERS parameter, as described in Chapter 9 of the *Net8 Administrator's Guide*.

Java, SQL, or PL/SQL clients, which execute Java stored procedures on the server, connect to the database over a Net8 connection. For a full description of how to configure the Net8 connection, see the *Net8 Administrator's Guide*.

Enterprise JavaBeans and CORBA Configuration

Clients access EJB and CORBA applications in the database over an Inter-Orb Protocol (IIOP) connection. IIOP is an implementation of GIOP over TCP/IP. To support an IIOP connection, you must configure the database in MTS mode with the General Inter-Orb Protocol (GIOP) presentation.

Oracle8i also supports the use of authentication data such as certificates and private keys required for use by SSL in combination with both types of GIOP protocols—regular GIOP and session-based GIOP.

For a complete description of how to configure MTS, see the *Net8 Administrator's Guide*. For a full description of how to configure the GIOP presentation, see the *Oracle8i Enterprise JavaBeans Developer's Guide and Reference* or the *Oracle8i CORBA Developer's Guide and Reference* books..

Enabling the Java Client

In order to run Java between the client and server, your client system must perform the following:

1. [Install JDK on the Client.](#)
2. [Set up CLASSPATH.](#)
3. [Verify the Port/SID.](#)
4. [Test Install with Samples.](#)

1. Install JDK on the Client

The client system is defined as the system where you execute the JServer tools, such as `loadjava`, `deployejb`. You can use the same system for both your client and server.

The client system requires JDK 1.1.6 or later. Solaris 2.6 bundles JDK 1.1.3, which does not work with our samples. Verify that your PATH includes JDK 1.1.6 or later and does not include JDK 1.1.3. To confirm what version of the JDK you are using, perform the following:

```
$ which java
/usr/local/packages/jdk1.1.6/bin/java
$ which javac
/usr/local/packages/jdk1.1.6/bin/javac
$ java -version
java version "1.1.6"
```

If JDK 1.1.6 does not appear within these commands, either put your JDK 1.1.6 installation at the start of PATH or remove the 1.1.3 installation. In addition, check your CLASSPATH for references to the incorrect JDK version.

Note: All Oracle8i Java-based client tools work in the Java 2 environment.

2. Set up CLASSPATH

If your client is a Java client involved with a distributed application—CORBA, EJB, or RMI—you must perform one of the following before compiling your client code:

- Set up CLASSPATH to include support JAR or ZIP files.
- Include support JAR or ZIP files within an option on the compile line.

For the Java client to work across nodes in a distributed application, it must be compiled with appropriate server stubs. The required JAR or ZIP files that must be included in the client's CLASSPATH are as follows:

Note: For NT users, the environment variables would be %ORACLE_HOME% and %JAVA_HOME%.

Basic Included JAR files:

```
$ORACLE_HOME/lib/aurora_client.jar
$ORACLE_HOME/lib/mts.jar
$ORACLE_HOME/lib/vbjorb.jar
$ORACLE_HOME/lib/vbjapp.jar
```

For execution, include `$ORACLE_HOME/lib/aurora.zip` on the CLASSPATH.

JAR Files Necessary for JDK 1.1 Clients

In order for a JDK 1.1 client to communicate with the Java 2 server, you must include the following JVM JAR file:

```
$JAVA_HOME/lib/classes.zip
```

For any interaction with JDBC, include the following ZIP file:

```
$ORACLE_HOME/jdbc/lib/classes111.zip
```

For any client that uses SSL, include the following JAR files:

```
$ORACLE_HOME/jlib/jssl-1_1.jar  
$ORACLE_HOME/jlib/javax-ssl-1_1.jar
```

JAR Files Necessary for Java 2 Clients

In order for a Java 2 client to communicate with the Java 2 server, you must make sure that one of the following JVM JAR files are in the CLASSPATH:

- For JDK 1.2, include `dt.jar`
- For JRE 1.2, include `rt.jar`

For any interaction with JDBC, include the following ZIP file:

```
$ORACLE_HOME/jdbc/lib/classes12.zip
```

For any client that uses SSL, include the following JAR files:

```
$ORACLE_HOME/jlib/jssl-1_2.jar  
$ORACLE_HOME/jlib/javax-ssl-1_2.jar
```

JAR Files Included for Clients that use SQLJ

```
$ORACLE_HOME/sqlj/lib/translator.zip
```

In addition to this files, add the appropriate `runtimeX.zip` file, as follows:

- Java 2 client using JDBC 8.1.7—`$ORACLE_HOME/sqlj/lib/runtime12.zip`
- Java 2 Enterprise Edition client using JDBC 8.1.7—`$ORACLE_HOME/sqlj/lib/runtime12ee.zip`
- JDK 1.1 client using JDBC 8.1.7—`$ORACLE_HOME/sqlj/lib/runtime11.zip`

- Any JDK client using JDBC 8.1.6 or previous version—`$ORACLE_HOME/sqlj/lib/runtime.zip`

JAR Files Included for Clients that use JSP

```
$ORACLE_HOME/jsp/lib/ojsp.jar
$ORACLE_HOME/lib/xmlparserv2.jar
$ORACLE_HOME/lib/servlet.jar
```

If your JSP pages use OracleJML tags or database utility JavaBeans, add the following:

```
$ORACLE_HOME/jsp/lib/ojsputil.jar
$ORACLE_HOME/jsp/lib/oraclexmlsql.jar
```

See the appropriate feature documentation for more information. You can also see examples of these JAR and ZIP files used in both the readme or makefile within the feature demo samples.

3. Verify the Port/SID

If you do not configure the default listener port numbers or database SID in your installation—such as ports 1521, 2481, and SID `orcl`—the samples will not work correctly. All the samples expect the default port numbers and database SID provided by an Oracle8i Typical install. If you have different values, specify the new port numbers as follows:

Host type	Directions
UNIX	<code>\$ make SERVICE=ess_iiop://localhost:myportnum:mysid</code>
Windows NT	In the control panel, set the system environment variable <code>ORACLE_SERVICE</code> to <code>ess_iiop://localhost:myportnum:mysid</code>

4. Test Install with Samples

We provide a set of samples in `$ORACLE_HOME/javavm/demo/demo.tar` (or `demo.zip` for Windows NT). These samples compile and run for a database installed with the Oracle8i Typical install option. Execute these samples as a test of your installation.

```
$ORACLE_HOME/javavm/demo/examples/jsp/helloworld
$ORACLE_HOME/javavm/demo/examples/corba/basic/helloworld
$ORACLE_HOME/javavm/demo/examples/ejb/basic/helloworld
```

If these samples do not compile or run, your environment is incorrect. Similarly, if these samples compile and run, but your code does not, then a problem exists within your build environment or code.

Note: It is important that you run these examples using the supplied Makefiles (or batch files on NT) when verifying your installation.

Verify that the samples work before using more complex build environments, such as Visual Cafe, JDeveloper, or VisualAge.

Security For Oracle8i Java Applications

Security is a large arena that includes network security for the connection, access and execution control of operating system resources or of JVM and user-defined classes, and bytecode verification of imported JAR files from an external source. You should be aware of what type of security you desire for your Java applications. The following sections describe the various security support available for Java applications within Oracle8i.

- [Network Connection Security](#)
- [Database Contents and JVM Security](#)

Network Connection Security

There are two major aspects to network security: authentication and data confidentiality. The type of authentication and data confidentiality is dependent on how you connect to the database—through Net8, JDBC, or distributed object (EJB or CORBA) connection.

Connection Security	Description
Net8	<p>The database can require both authentication and authorization before allowing a user to connect to the database. Net8 database connection security can require one or more of the following:</p> <ul style="list-style-type: none">■ Use a username and password for client verification. Each incoming connection into the database has to provide the correct username/password configured within Net8. For more information, see the <i>Net8 Administrator's Guide</i>.■ Use Advanced Networking Option for encryption, kerberos, or secureId. See the <i>Oracle Advanced Security Administrator's Guide</i>.■ Use SSL for certificate authentication. See the <i>Oracle Advanced Security Administrator's Guide</i>.
JDBC	<p>JDBC connection security required is similar to the constraints required on a Net8 database connection. In addition to the books listed in the Net8 database connection section, see the <i>Oracle8i JDBC Developer's Guide and Reference</i>.</p>
Distributed Object	<p>Encryption and authentication might be required for distributed applications, such as EJB and CORBA. For more information, see the <i>Oracle8i Enterprise JavaBeans Developer's Guide and Reference</i> or the <i>Oracle8i CORBA Developer's Guide and Reference</i> books.</p>

Database Contents and JVM Security

Once you are connected to the database, you still must have the correct Java 2 Permissions and database privileges to access the resources stored within the database. These resources include the following:

- database resources, such as tables and PL/SQL packages
- operating system resources, such as files and sockets
- Aurora JVM classes
- user-loaded classes

These resources can be protected by the following two methods:

Resource Security	Description
Database Resource Security	<p>Authorization for database resources requires that database privileges (not the same as the Java 2 security permissions) are granted to resources. For example, database resources include tables, classes, or PL/SQL packages. For more information, see the <i>Oracle8i Application Developer's Guide - Fundamentals</i>.</p> <p>All user-defined classes are secured against users from other schemas. You can grant execution permission to other users/schemas through an option on the <code>loadjava</code> command. For more information on setting execution rights when loading the class, see the <code>-grant</code> option discussed in "Loading Classes" on page 2-23 or the <i>Oracle8i Java Tools Reference</i> for complete information on <code>loadjava</code>.</p>
JVM Security	<p>JServer uses Java 2 security, which uses Permission objects to protect operating system resources. Java 2 security is automatically installed upon startup and protects all operating system resources and JVM classes from all users, except <code>JAVA_ADMIN</code>. <code>JAVA_ADMIN</code> can grant permission to other users to access these classes.</p> <p>See "Java 2 Security" on page 5-3 for how to manage and modify Java 2 Permissions and policies.</p>

Java 2 Security

Each user or schema must be assigned the proper permissions to access operating system resources. For example, this includes sockets, files, and system properties.

Java 2 security was created to provide a flexible, configurable security for Java applications. With Java 2 security, you can define exactly what permissions on each loaded object a schema or role will have. In 8.1.5, the security provided you the choice of two secure roles:

- `JAVAUSERPRIV`—few Permissions, including examining properties
- `JAVASYSPRIV`—major Permissions, including updating JVM protected packages

Note: Both roles still exist within 8.1.6 for backward compatibility; however, Oracle recommends that you specify each Permission explicitly, rather than utilize these roles.

Because JServer security is based on Java 2 security, you assign Permissions on a class by class basis. Permissions contains two string attributes:

- target (name) attribute
- action attribute

These permissions are assigned through database management tools. Each permission is encapsulated in a Permission object and is stored within a Permission table. The methods for managing all permissions are the subject for most of this chapter.

Java security was created for the non-database world. When you apply the Java 2 security model within the database, certain differences manifest themselves. For example, Java 2 security defines that all applets are implicitly untrusted and all classes within the CLASSPATH are trusted. Within Oracle8i, all classes are loaded within a secure database; thus, no classes are trusted.

The following table briefly describes the differences between Sun Microsystem's Java 2 security and Oracle8i's implementation. This table assumes that you already understand Sun Microsystem's Java 2 security model. For more information, we recommend the following books:

- *Inside Java 2 Platform Security* by Li Gong
- *Java Security* by Scott Oaks

Java 2 Security standard	Oracle8i implementation
Java classes located within the CLASSPATH are trusted.	All Java classes are loaded within the database. Classes are trusted on a class by class basis according to the Permission granted.
You can specify the policy through the <code>-usepolicy</code> flag on the <code>java</code> command line.	You must specify the policy within the <code>PolicyTable</code> .

Java 2 Security standard	Oracle8i implementation
You can write your own SecurityManager or use the Launcher.	You can write your own SecurityManager; Oracle recommends that you use only Aurora's SecurityManager or that you extend Aurora's SecurityManager. If you want to modify the behavior, you should not define a SecurityManager; instead, you should extend <code>oracle.aurora.rdbms.SecurityManagerImpl</code> and override specific methods.
SecurityManager is not initialized for you. You must initialize the SecurityManager.	Aurora always initializes SecurityManager at startup.
Permissions are determined by the location where the application or applet is loaded from (the URL) or keycode (signed code).	Permissions are determined by the schema in which the class is loaded. JServer does not support signed code.
The security policy is defined in a file.	The PolicyTable definition is contained within a secure database table.
You can update the security policy file through a text editor (if you have the correct Permissions) or through a tool.	You update the PolicyTable through DBMS_JAVA procedures. After initialization, only JAVA_ADMIN has permission to modify the PolicyTable. JAVA_ADMIN must grant you the right to modify the PolicyTable for you to grant Permissions to others.
Permissions are assigned to a protection domain, which classes can belong to.	All classes within the same schema are within the same protection domain.
Use <code>CodeSource</code> class for identifying code.	Use <code>CodeSource</code> class for identifying schema.
<ul style="list-style-type: none"> ■ The <code>equals</code> method returns true if the URL and certificates are equal. ■ The <code>implies</code> method returns true if the first <code>CodeSource</code> is a generic representation that includes the specific <code>CodeSource</code> object. 	<ul style="list-style-type: none"> ■ The <code>equals</code> method returns true if the schemas are the same. ■ The <code>implies</code> method returns true if the schemas are the same.
Supports positive Permissions only (grant).	Supports both positive (grant) and limitation (restrict) Permissions.

Setting Permissions

As designed in Java 2 security, Oracle8i supports the security classes. Normally, you set the Permissions for the code base either through a tool or by editing the security policy file. In Oracle8i, you set the Permissions dynamically through DBMS_JAVA procedures. These procedures modify a policy table, which is a new table within the database that exclusively manages Java 2 security Permissions.

Two views have been created for you to view the policy table: USER_JAVA_POLICY and DBA_JAVA_POLICY. Both views contain information about granted and limitation Permissions. The DBA_JAVA_POLICY view can see all rows within the policy table; the USER_JAVA_POLICY table can only see Permissions relevant to the current user. The following is a description of the rows within each view:

Table Column	Description
Kind	GRANT or RESTRICT. Shows whether this Permission is a positive (GRANT) or a limitation (RESTRICT) Permission.
Grantee	The name of the user, schema, or role that the Permission object is assigned to.
Permission_schema	The schema that the Permission object is loaded into.
Permission_type	The Permission class type, which is designated by a string containing the full class name, such as, <code>java.io.FilePermission</code> .
Permission_name	The target attribute (name) of the Permission object. You use this name when defining the Permission. When defining the target for a Permission of type <code>PolicyTablePermission</code> , the name can become quite complicated. See " Acquiring Administrative Permission to Update Policy Table " on page 5-12 for more information.
Permission_action	The action attribute for this Permission. Many Permissions expect a null value if no action is appropriate for the Permission.
Status	ACTIVE or INACTIVE. After creating a row for a Permission, you can disable or re-enable it. This column shows the status of whether the Permission is enabled (ACTIVE) or disabled (INACTIVE).
Key	Sequence number you use to identify this row. This number should be supplied when disabling, enabling, or deleting the Permission.

There are two ways to set your Permissions:

- **Fine-Grain Definition for Each Permission**—You grant each Permission individually for specific users or roles. If you do not grant a Permission for access, the schema will be denied access.

- **General Permission Definition Assigned to Roles**—If you do not want to grant specific Permissions for each user, you can grant roles, which grants a collection of Permissions to the user. Oracle8i supplies the roles: JAVAUSERPRIV or JAVASYSPRIV.

Note: For absolute certainty about your security, you should implement the fine-grain definition. The general definition is easier; but you might not get the exact security you require.

Fine-Grain Definition for Each Permission

To set individual Permissions within the policy table, you must provide the following information:

Parameter	Description
Grantee	The name of the user, schema, or role that you want the grant to apply to. PUBLIC specifies that the row applies to all users.
Permission type	The Permission class that you are granting Permission on. For example, if you are defining access to a file, the Permission type would be <code>FilePermission</code> . This parameter requires a fully-qualified name of a class that extends <code>java.lang.security.Permission</code> . If the class is not within SYS, the name should be prefixed by <code><schema>.</code> For example, <code>mySchema:myPackage.MyPermission</code> is a valid name for a user generated Permission.
Permission name	The meaning of the target attribute is defined by the Permission class. Examine the appropriate Permission class for the relevant name.
Permission action	The type of actions that you can specify vary according to the Permission type. For example, <code>FilePermission</code> can have the actions of read or write.
Key	Number returned from grant or limit to use on enable, disable, or delete methods.

You can grant either Java 2 Permissions or create your own. The Java 2 Permissions are listed in [Table 5-1](#). If you would like to create your own Permissions, see "[Creating Permissions](#)" on page 5-14.

Table 5–1 Permission Types

- `java.util.PropertyPermission`
 - `java.io.SerializablePermission`
 - `java.io.FilePermission`
 - `java.net.NetPermission`
 - `java.net.SocketPermission`
 - `java.lang.RuntimePermission`
 - `java.lang.reflect.ReflectPermission`
 - `java.security.SecurityPermission`
 - `oracle.aurora.rdbms.security.PolicyTablePermission`
 - `oracle.aurora.security.JServerPermission`
-

You can grant permissions using either SQL or Java, as shown below. However, each returns a row key identifier that identifies the row within the permission table. In the Java version of `DBMS_JAVA`, each method returns the row key identifier, either as a returned parameter or as an OUT variable in the parameter list. In the PL/SQL `DBMS_JAVA` package, the row key is returned only in the procedure that defines the key OUT parameter. This key is used to enable and disable specific Permissions. See "[Enabling or Disabling Permissions](#)" on page 5-17 for more information.

If, after executing the grant, a row already exists for the exact Permission, no update occurs, but the key for that row is returned. If the row was disabled, executing the grant enables the existing row.

Note: If granting `FilePermission`, you must provide the physical name of the directory or file, such as `/private/oracle`. You cannot provide either an environment variable, such as `$ORACLE_HOME`, or a symbolic link. Also, to denote all files within a directory, provide the `*` symbol, as follows: `'/private/oracle/*'`. To denote all directories and files within a directory, provide the `-` symbol, as follows: `'/private/oracle/-'`.

Granting Permissions using the `DBMS_JAVA` package:


```
procedure grant_permission( grantee varchar2, permission_type varchar2,
                             permission_name varchar2,
                             permission_action varchar2 )
```

```
procedure grant_permission( grantee varchar2, permission_type varchar2,
                             permission_name varchar2,
                             permission_action varchar2, key OUT number)
```

Granting Permissions using Java:

```
long oracle.aurora.rdbms.security.PolicyTableManager.grant(
    java.lang.String grantee,
    java.lang.String permission_type,
    java.lang.String permission_name,
    java.lang.String permission_action);
```

```
void oracle.aurora.rdbms.security.PolicyTableManager.grant(
    java.lang.String grantee,
    java.lang.String permission_type,
    java.lang.String permission_name,
    java.lang.String permission_action,
    long[] key);
```

Limiting Permissions using the DBMS_JAVA package:

```
procedure restrict_permission( grantee varchar2, permission_type varchar2,
                                permission_name varchar2,
                                permission_action varchar2)
```

```
procedure restrict_permission( grantee varchar2, permission_type varchar2,
                                permission_name varchar2,
                                permission_action varchar2, key OUT number)
```

Limiting Permissions using Java:

```
long oracle.aurora.rdbms.security.PolicyTableManager.restrict(
    java.lang.String grantee,
    java.lang.String permission_type,
    java.lang.String permission_name,
    java.lang.String permission_action);
```

```
void oracle.aurora.rdbms.security.PolicyTableManager.restrict(
    java.lang.String grantee,
    java.lang.String permission_type,
    java.lang.String permission_name,
```

```
java.lang.String permission_action,  
long[] key);
```

Example 5-1 Granting Permissions

Assuming that you have appropriate Permissions to modify the policy table, you use the `grant_permission` method within the `DBMS_JAVA` package to modify the `PolicyTable` to allow the user access to the indicated file. In this example, the user, Larry, has `PolicyTable` modification Permission. Within a SQL package, Larry grants permission to read and write a file to the user Dave.

```
connect larry/larry
```

```
REM Grant DAVE permission to read and write the Test1 file.  
call dbms_java.grant_permission('DAVE',  
                                'java.io.FilePermission', '/test/Test1',  
                                'read,write');
```

```
REM commit the changes to the PolicyTable  
commit;
```

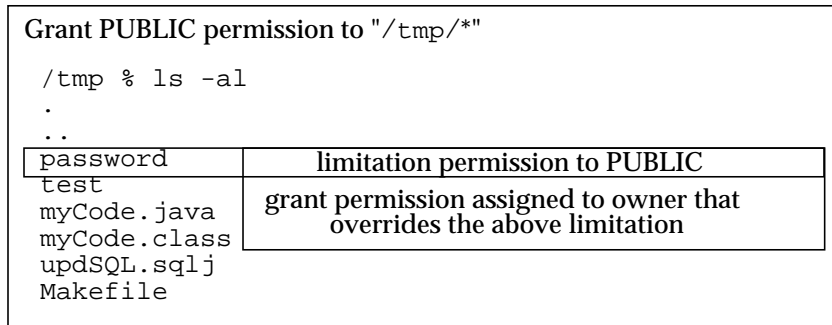
Example 5-2 Limiting Permissions

You use the `restrict` method for specifying a limitation or exception for general rules. A general rule is a rule where, in most cases, the Permission is true. However, there may be exceptions to this rule. For these exceptions, you specify a limitation Permission.

That is, if you have defined a general rule that no one can read or write for an entire directory, you can define a limitation on an aspect of this rule through the `restrict` method. For example, if you want to allow access to all files within the `/tmp` directory—except for your `password` file that exists in that directory—you would grant permission for read and write to all files within `/tmp` and limit read and write access to the `password` file.

If you want to specify an exception to the limitation, you would create an explicit grant Permission to override the limitation Permission. In the scenario mentioned above, if you want the file owner to still be able to modify the `password` file, you can grant a more explicit Permission to allow access to one user, which will override

the limitation. JServer security combines all rules to understand who really has access to the `password` file. This is demonstrated in the following diagram:



The explicit rule is as follows:

If the limitation Permission implies the request, then for a grant to be effective, the limitation Permission must also imply the grant.

The following is the code that implements this example:

1. Grant everyone (PUBLIC) read and write permission to all files in `/tmp`.
2. Limit everyone (PUBLIC) from reading or writing one file in `/tmp/password`.
3. Grant Larry (owner) explicit permission to read and write the `password` file.

```
connect larry/larry
```

```
REM Grant permission to all users (PUBLIC) to be able to read and write
REM all files in /tmp.
```

```
call dbms_java.grant_permission('PUBLIC',
                                'java.io.FilePermission',
                                '/tmp/*',
                                'read,write');
```

```
REM Limit permission to all users (PUBLIC) from reading or writing the
REM password file in /tmp.
```

```
call dbms_java.restrict_permission('PUBLIC',
                                    'java.io.FilePermission',
                                    '/tmp/password',
                                    'read,write');
```

```
REM By providing a more specific rule that overrides the limitation,
REM Larry can read and write /tmp/password.
```

```
call dbms_java.grant_permission('LARRY',
                                'java.io.FilePermission',
                                '/tmp/password',
                                'read,write');

commit;
```

Acquiring Administrative Permission to Update Policy Table

All Permissions are rows within the policy table. As it is a table within the database and thus a resource, permission is needed to modify it. Specifically, the `PolicyTablePermission` object is required to modify the table. After the first initialization for JServer, only a single role—`JAVA_ADMIN`—is granted the `PolicyTablePermission` to modify the policy table. The `JAVA_ADMIN` role is immediately assigned to `DBA`; thus, if you are assigned to the `DBA` group, you will automatically take on all `JAVA_ADMIN` Permissions.

In order for you to be able to add Permissions as rows to this table, `JAVA_ADMIN` must grant your schema update rights for the `PolicyTablePermission`. This Permission defines that your schema can add rows to the table. Each `PolicyTablePermission` is for a specific Permission type. For example, in order for you to add a Permission that controls access to a file, you must have a `PolicyTablePermission` that allows you to grant or limit a Permission on a `FilePermission`. Once this occurs, you have administrative Permission for `FilePermission`.

The administrator could grant and limit the `PolicyTablePermissions` in the same manner as other Permissions, but the syntax is complicated. For ease of use, use one of the following methods within the `DBMS_JAVA` package to grant administrative Permissions.

Granting policy table administrative Permissions using `DBMS_JAVA`:

```
procedure grant_policy_permission( grantee varchar2, permission_schema varchar2,
                                  permission_type varchar2,
                                  permission_name varchar2)
```

```
procedure grant_policy_permission( grantee varchar2, permission_schema varchar2,
                                  permission_type varchar2,
                                  permission_name varchar2,
                                  key OUT number)
```

Granting policy table administrative permission using Java:

```
long oracle.aurora.rdbms.security.PolicyTableManager.grantPolicyPermission(
    java.lang.String grantee,
    java.lang.String permission_type,
```

```
java.lang.String permission_name);
```

```
void oracle.aurora.rdbms.security.PolicyTableManager.grantPolicyPermission(
    java.lang.String grantee,
    java.lang.String permission_type,
    java.lang.String permission_name,
    long[] key);
```

Parameter	Description
Grantee	The name of the user, schema, or role that you want the grant to apply to. PUBLIC specifies that the row applies to all users.
Permission_schema	The <schema> where the Permission class is loaded.
Permission_type	The Permission class that you are granting Permission on. For example, if you are defining access to a file, the Permission type would be FilePermission. This parameter requires a fully-qualified name of a class that extends java.lang.security.Permission. If the class is not within SYS, the name should be prefixed by <schema>:. For example, mySchema:myPackage.MyPermission is a valid name for a user generated Permission.
Permission_name	The meaning of the target attribute is defined by the Permission class. Examine the appropriate Permission class for the relevant name.
Row_number	Number returned from grant or limitation to use on enable, disable, or delete methods.

Note: When looking at the policy table, the name within the PolicyTablePermission rows contains both the Permission type and the Permission name, which are separated by a '#'. For example, to grant a user administrative rights for reading a file, the name in the row contains java.io.FilePermission#read. The '#' separates the Permission class from the Permission name.

Example 5-3 Granting PolicyTable Permission

The following example shows JAVA_ADMIN (as SYS) giving Larry permission to update the PolicyTable for FilePermission. Once this Permission is granted, Larry can grant permissions to other users for reading, writing, and deleting files.

```
REM Connect as SYS, which is assigned JAVA_ADMIN role, to give Larry permission
REM to modify the PolicyTable
connect SYS/SYS as SYSDBA
```

```
REM SYS grants Larry the right to administer permissions for
REM FilePermission
call dbms_java.grant_policy_permission('LARRY', 'SYS',
                                     'java.io.FilePermission', '*');
```

Creating Permissions

You can create your own Permission type by performing the following:

1. [Create and load the user Permission.](#)
2. [Grant administrative and action Permissions to specified users.](#)
3. [Implement security checks using the Permission.](#)

1. Create and load the user Permission You can create your own Permission by extending the Java 2 Permission class. Any user created Permission must extend Permission. The following example creates MyPermission. MyPermission extends BasicPermission, which in turn extends Permission.

```
package test.larry;
import java.security.Permission;
import java.security.BasicPermission;

public class MyPermission extends BasicPermission {

    public MyPermission(String name) {
        super(name);
    }

    public boolean implies(Permission p) {
        boolean result = super.implies(p);
        return result;
    }
}
```

2. Grant administrative and action Permissions to specified users When you create a Permission, you are designated as owner of that Permission. The owner is implicitly granted administrative Permission. This means that the owner can be an administrator for this Permission and can execute `grant_policy_permission`. Administrative Permission permits the user to update the policy table for the user-defined Permission.

For example, if LARRY creates a Permission, MyPermission, only LARRY can invoke `grant_policy_permission` for himself or another user. This method updates the PolicyTable on who can grant rights to MyPermission. The following code demonstrates this:

```
REM Since Larry is the user that owns MyPermission, Larry connects to
REW the database to assign permissions for MyPermission.
connect larry/larry
```

```
REM As the owner of MyPermission, Larry grants himself the right to
REM administer permissions for test.larry.MyPermission within the JVM
REM security PolicyTable. Only the owner of the user-defined permission
REM can grant administrative rights.
call dbms_java.grant_policy_permission('LARRY', 'LARRY',
                                     'test.larry.MyPermission', '*');
```

```
REM commit the changes to the PolicyTable
commit;
```

Once you have granted administrative rights, you can grant action Permissions for the user created Permission. For example, the following SQL grants permission for LARRY to execute anything within MyPermission and DAVE to be able to only execute actions that start with "act".

```
REM Since Larry is the user that creates MyPermission, Larry connects to
REW the database to assign permissions for MyPermission.
connect larry/larry
```

```
REM Once able to modify the PolicyTable for MyPermission, Larry grants himself
REM full permission for MyPermission. Notice that the Permission is prepended
REM with its owner schema.
call dbms_java.grant_permission( 'LARRY',
                                'LARRY:test.larry.MyPermission', '*', null);
```

```
REM Larry grants Dave permission to do any actions that start with 'act.*'.
call dbms_java.grant_permission
    ('DAVE', 'LARRY:test.larry.MyPermission', 'act.*', null);
```

```
REM commit the changes to the PolicyTable
commit;
```

3. Implement security checks using the Permission Once you have created, loaded, and assigned Permissions for MyPermission, you must implement the call to SecurityManager for having the Permission checked. There are four methods in the following example: sensitive, act, print, and hello. Because of the

Permissions granted in the SQL example in step 2, the following users can execute methods within the example class:

- LARRY can execute any of the methods.
- DAVE is given permission to execute only the `act` method.
- Anyone can execute the `print` and `hello` methods. The `print` method does not check any Permissions, so anyone can execute the `print` method. The `hello` method executes `AccessController.doPrivileged`, which means that the method executes with LARRY's Permissions. This is referred to as **definer's rights**.

```
package test.larry;
import java.security.AccessController;
import java.security.Permission;
import java.security.PrivilegedAction;

import java.sql.Connection;
import java.sql.SQLException;

/**
 * MyActions is a class with a variety of public methods that
 * have some security risks associated with them. We will rely
 * on the Java security mechanisms to ensure that they are
 * performed only by code that is authorized to do so.
 */

public class Larry {

    private static String secret = "Larry's secret";

    MyPermission sensitivePermission = new MyPermission("sensitive");

    /**
     * This is a security sensitive operation. That is it can
     * compromise our security if it is executed by a "bad guy".
     * Only larry has permission to execute sensitive.
     */
    public void sensitive() {
        checkPermission(sensitivePermission);
        print();
    }

    /**
     * Will print a message from Larry. We need to be
```



```

    * careful about who is allowed to do this
    * because messages from Larry may have extra impact.
    * Both larry and dave have permission to execute act.
    */
public void act(String message) {
    MyPermission p = new MyPermission("act." + message);
    checkPermission(p);
    System.out.println("Larry says: " + message);
}

/**
 * Print our secret key
 * No permission check is made; anyone can execute print.
 */
private void print() {
    System.out.println(secret);
}

/**
 * Print "Hello"
 * This method invokes doPrivileged, which makes the method run
 * under definer's rights. So, this method runs under Larry's
 * rights, so anyone can execute hello.
 * Only Larry can execute hello
 */
public void hello() {
    AccessController.doPrivileged(new PrivilegedAction() {
        public Object run() { act("hello"); return null; }
    });
}

/**
 * If a security manager is installed ask it to check permission
 * otherwise use the AccessController directly
 */
void checkPermission(Permission permission) {
    SecurityManager sm = System.getSecurityManager();
    sm.checkPermission(permission);
}
}

```

Enabling or Disabling Permissions

Once you have created a row that defines a Permission, you can disable it so that it is no longer applied. However, if you decide you want the row action again, you

can enable the row. You can delete the row from the table if you believe that it will never be used again. To delete, you must first disable the row. If you do not disable the row, the deletion will not occur.

To disable rows, you can use either the `disable_permission` or the `revoke` method.

- The `revoke_permission` method takes in parameters similar to the `grant` and `restrict` methods. It searches the entire policy table for all rows that match the supplied parameters.
- The `disable_permission` method disables only a single row within the policy table. To do this, it takes in the policy table key. This key is also necessary to enable or delete a Permission. To retrieve the Permission key number, perform one of the following:
 - * Save the key when it is returned on the `grant` or `limit` calls. If you do not foresee a need to ever enable or disable the Permission, you can use the `grant` and `limit` calls that do not return the Permission number.
 - * View `DBA_JAVA_POLICY` or `USER_JAVA_POLICY` for the appropriate Permission key number.

Disabling Permissions using `DBMS_JAVA`:

```
procedure revoke_permission(permission_schema varchar2,  
                             permission_type varchar2,  
                             permission_name varchar2,  
                             permission_action varchar2)
```

```
procedure disable_permission(key number)
```

Disabling Permissions using Java:

```
void revoke(String schema, String type, String name, String action);  
  
void oracle.aurora.rdbms.security.PolicyTableManager.disable(long number);
```

Enabling Permissions using `DBMS_JAVA`:

```
procedure enable_permission(key number)
```

Enabling Permissions using Java:

```
void oracle.aurora.rdbms.security.PolicyTableManager.enable(long number);
```

Deleting Permissions using `DBMS_JAVA`:

procedure **delete_permission**(key number)

Deleting Permissions using Java:

```
void oracle.aurora.rdbms.security.PolicyTableManager.delete(long number);
```

Permission Types

[Table 5–2](#) lists the installed Permission types. Whenever you want to grant or limit a Permission, you must provide the Permission type within the DBMS_JAVA method. The Permission types with which you control access are the following:

- Oracle-provided Permission types listed in [Table 5–2](#)
- user created Permission types that extend `java.security.Permission`

Table 5–2 Permission Types

-
- `java.util.PropertyPermission`
 - `java.io.SerializablePermission`
 - `java.io.FilePermission`
 - `java.net.NetPermission`
 - `java.net.SocketPermission`
 - `java.lang.RuntimePermission`
 - `java.lang.reflect.ReflectPermission`
 - `java.security.SecurityPermission`
 - `oracle.aurora.rdbms.security.PolicyTablePermission`
 - `oracle.aurora.security.JServerPermission`
-

All the Java Permission types are documented in Sun Microsystem's Java 2 documentation.

Note: SYS is granted permission to load libraries that come with Oracle. However, Aurora does not support other users loading libraries, because loading C within the database is insecure. So, you are not allowed to grant permission for `loadLibrary.*` of `RuntimePermission`.

The Oracle-specific Permissions, `PolicyTablePermission` and `JServerPermission` are described below:

oracle.aurora.rdbms.security.PolicyTablePermission This Permission controls who can update the policy table. Once granted the right to update the policy table for a certain `Permission` type, the user can control other user's access to some resource.

After JServer initialization, only the `JAVA_ADMIN` role can grant administrative rights for the policy table through `PolicyTablePermission`. Once it grants this right to other users, these users can in turn update the policy table with their own grant and limitation Permissions.

To grant policy table updates, you use the `DBMS_JAVA` method: `grant_policy_permission`, as discussed in "[Acquiring Administrative Permission to Update Policy Table](#)" on page 5-12. Once you have updated the table, you can view either the `DBA_JAVA_POLICY` or `USER_JAVA_POLICY` views to see who has been granted Permissions.

oracle.aurora.security.JServerPermission You use this Permission to grant and limit access to Aurora JVM resources. The `JServerPermission` extends from `BasicPermission`. The following table lists the names that `JServerPermission` grants access for:

Permission Name	Description
<code>LoadClassInPackage.<package_name></code>	grants the ability to load a class within the specified package
<code>Verifier</code>	grants the ability to turn the bytecode verifier on or off
<code>Debug</code>	grants the ability for debuggers to connect to a session
<code>JRIExtensions</code>	grants the use of <code>MEMSTAT</code>
<code>Memory.Call</code>	grants rights to call certain methods in <code>oracle.aurora.vm.OracleRuntime</code> on call settings
<code>Memory.Stack</code>	grants rights to call certain methods in <code>oracle.aurora.vm.OracleRuntime</code> on stack settings
<code>Memory.SGAIntern</code>	grants rights to call certain methods in <code>oracle.aurora.vm.OracleRuntime</code> on SGA settings

Permission Name	Description
Memory.GC	grants rights to call certain methods in <code>oracle.aurora.vm.OracleRuntime</code> on garbage collector settings

Initial Permission Grants

When you first initialize JServer, several roles are populated with certain Permission grants. The following tables show these roles and their initial Permissions:

1. The `JAVA_ADMIN` role is given access to modify the policy table for all Permissions. All DBAs, including `SYS`, are granted `JAVA_ADMIN`. Full administrative rights to update the policy table are granted for the following Permissions:

-
- `java.util.PropertyPermission`
 - `java.io.SerializablePermission`
 - `java.io.FilePermission`
 - `java.net.NetPermission`
 - `java.net.SocketPermission`
 - `java.lang.RuntimePermission`
 - `java.lang.reflect.ReflectPermission`
 - `java.security.SecurityPermission`
 - `oracle.aurora.rdbms.security.PolicyTablePermission`
 - `oracle.aurora.security.JServerPermission`
-

2. In addition to the `JAVA_ADMIN` Permissions, `SYS` is also granted the following Permissions:

Note: Within the `RuntimePermission` grants, there seems to be unnecessary granting of more specific `Permission` for `loadLibrary.<package>`. The reason for this is to override the limitation given to `PUBLIC` for `loadLibrary.*`.

Table 5–3 *SYS Initial Permissions*

Permission Type	Permission Name	Action Granted
<code>oracle.aurora.rdbms.security. PolicyTablePermission</code>	*	Administrative rights to modify the policy table
<code>oracle.aurora.security.JServerPermission</code>	*	null
<code>java.net.NetPermission</code>	*	null
<code>java.security.SecurityPermission</code>	*	null
<code>java.util.PropertyPermission</code>	*	write
<code>java.lang.reflect.ReflectPermission</code>	*	null
<code>java.lang.RuntimePermission</code>	*	null
	<code>loadLibrary.xaNative</code>	null
	<code>loadLibrary.corejava</code>	null
	<code>loadLibrary.corejava_d</code>	null

- All users are initially granted the following Permissions. For the `JServerPermission`, all users can load classes, except for the list specified in the table. These exceptions are limitation Permissions. For more information on limitation Permissions, see [Example 5–2](#).

Table 5–4 *PUBLIC Default Permissions*

Permission Type	Permission Name	Granted Action
<code>oracle.aurora.rdbms.security. PolicyTablePermission</code>	<code>java.lang.RuntimePermission. loadLibrary.*</code>	null
<code>java.util.PropertyPermission</code>	*	read
	<code>user.language</code>	write

Table 5–4 PUBLIC Default Permissions

Permission Type	Permission Name	Granted Action
java.lang.RuntimePermission	exitVM	null
	createSecurityManager	null
	modifyThread	null
	modifyThreadGroup	null
	loadClassInPackage.* except for loadClassInPackage.java.*, loadClassInPackage.oracle.aurora.*, and loadClassInPackage.jdbc.*	null
oracle.aurora.security. JServerPermission		

Table 5–5 JAVAUSERPRIV Permissions

Permission Type	Permission Name	Action
java.net.SocketPermission	*	connect, resolve
java.io.FilePermission	<<ALL FILES>>	read
java.lang.RuntimePermission	modifyThreadGroup, stopThread, getProtectionDomain, readFileDescriptor, accessClassInPackage.*, and defineClassInPackage.*	null

Table 5–6 JAVASYSPRIV Permissions

Permission Type	Permission Name	Action
java.io.SerializablePermission	*	no applicable action
java.io.FilePermission	<<ALL FILES>>	read, write, execute, delete
java.net.SocketPermission	*	accept, connect, listen, resolve
java.lang.RuntimePermission	createClassLoader	null
	getClassLoader	null
	setContextClassLoader	null
	setFactory	null

Table 5–6 JAVASYSPRIV Permissions (Cont.)

Permission Type	Permission Name	Action
	setIO	null
	setFileDescriptor	null
	readFileDescriptor	null
	writeFileDescriptor	null

Table 5–7 JAVADEBUGPRIV Permissions

Permission Type	Permission Name	Action
oracle.aurora.security.JServerPermission	Debug	null
java.net.SocketPermission	*	connect, resolve

General Permission Definition Assigned to Roles

In 8.1.5, JVM security was controlled by granting the roles of JAVASYSPRIV, JAVAUSERPRIV, or JAVADEBUGPRIV to schemas. In the current version, these roles still exist as Permission groups. See the previous section, "[Initial Permission Grants](#)" on page 5-21 for the explicit Permissions set for each role. You can set up and define your own collection of Permissions. Once defined, you can grant any collection of Permissions to any user. That user will then have the same Permissions that exist within the role. In addition, if you need additional Permissions, you can add individual Permissions to either your specified user or role. Permissions defined within the policy table have a cumulative effect. See "[Fine-Grain Definition for Each Permission](#)" on page 5-7 for information on how to grant Permissions to a user or a role.

Note: The ability to write to properties, granted through the write action on `PropertyPermission`, is no longer granted to all users. Instead, you must either have `JAVA_ADMIN` grant you this Permission or you can receive this Permission by being granted the role of JAVASYSPRIV.

The following example gives Larry and Dave the following Permissions:

- Larry receives JAVASYSPRIV Permissions.

- Dave receives JAVADEBUGPRIV Permissions and the ability to read and write all files on the system.

```
REM Granting Larry the same permissions as exist within JAVASYSPRIV
grant javasyspriv to larry;
```

```
REM Granting Dave the ability to debug
grant javadebugpriv to dave;
```

```
commit;
```

```
REM I also want Dave to be able to read and write all files on the system
call dbms_java.grant_permission('DAVE', 'SYS:java.io.FilePermission',
                               '<<ALL FILES>>', 'read,write', null);
```

Debugging Permissions

A debug role, JAVADEBUGPRIV, was created to grant Permissions for running the debugger. The Permissions assigned to this role are listed in [Table 5-7](#). In order to have permission to invoke the debug agent, the caller must have been granted JAVADEBUGPRIV or the debug JServerPermission as follows:

```
REM Granting Dave the ability to debug
grant javadebugpriv to dave;
```

```
REM Larry grants himself permission to start the debug agent.
call dbms_java.grant_permission
    ('LARRY', 'oracle.aurora.security.JServerPermission', 'Debug', null);
```

A debugger provides extensive access to both code and data on the server, but at this time, we envision its use to be limited to development environments. Refer to the discussion in the section "[Debugging Server Applications](#)" on page 3-18 for information on using the debugging facilities in this release.

Permission for Loading Classes

In order to load classes, you must have the following Permission:

```
JServerPermission("LoadClassInPackage." + <class_name>)
```

The class name is the fully qualified name of the class that you are loading.

This excludes loading into system packages or replacing any system classes. Even if you are granted permission to load a system class, JServer prevents you from performing the load. System classes are classes that are installed by Oracle8i with

CREATE JAVA SYSTEM. The following error is thrown if you try to replace a system class:

```
ORA-01031 "Insufficient privileges"
```

The following shows the ability of each user after database installation, including Permissions and JServer restrictions:

- SYS can load any class, except for system classes.
- Any user can load classes in its own schema that do not start with the following patterns: `java.*`, `oracle.aurora.*`, `oracle.jdbc.*`. If the user wants to load such classes into another schema, it must be granted the `JServerPermission(LoadClassInPackage.<class>)` Permission.

The following example shows how to grant SCOTT Permission to load classes into the `oracle.aurora.*` package:

```
dbms_java.grant_permission('SCOTT', 'SYS:oracle.aurora.tools.*', null);
```

Oracle8i Java Application Performance

You can increase your Java application performance through one of the following methods:

- [Natively Compiled Code](#)
- [Java Memory Usage](#)
- [End-of-Call Migration](#)
- [Memory Profiling Utility](#)

Natively Compiled Code

The Java language was designed for a platform-independent, secure development model. To accomplish these goals, some execution performance was sacrificed. Translating Java bytecodes into machine instructions degrades performance. To regain some of the performance loss, you may choose to natively compile certain classes. For example, you may decide to natively compile code with CPU intensive classes.

Without native compilation, the Java code you load to the server is interpreted and the underlying core classes upon which your code relies (`java.lang.*`) are natively compiled.

Native compilation provides a speed increase ranging from two to ten times the speed of the bytecode interpretation. The exact speed increase is dependent on several factors, including:

- use of numerics
- degree of polymorphic message sends
- use of direct field access, as opposed to accessor methods
- amount of Array accessing
- casts

Because Java bytecodes were designed to be compact, natively compiled code can be considerably larger than the original bytecode. However, because the native code is stored in a shared library, it is shared among all users of the database.

Most JVMs use Just-In-Time compilers that convert the Java bytecodes to native machine instructions when methods are invoked. The JServer Accelerator uses an Ahead-Of-Time approach to recompiling the Java classes.

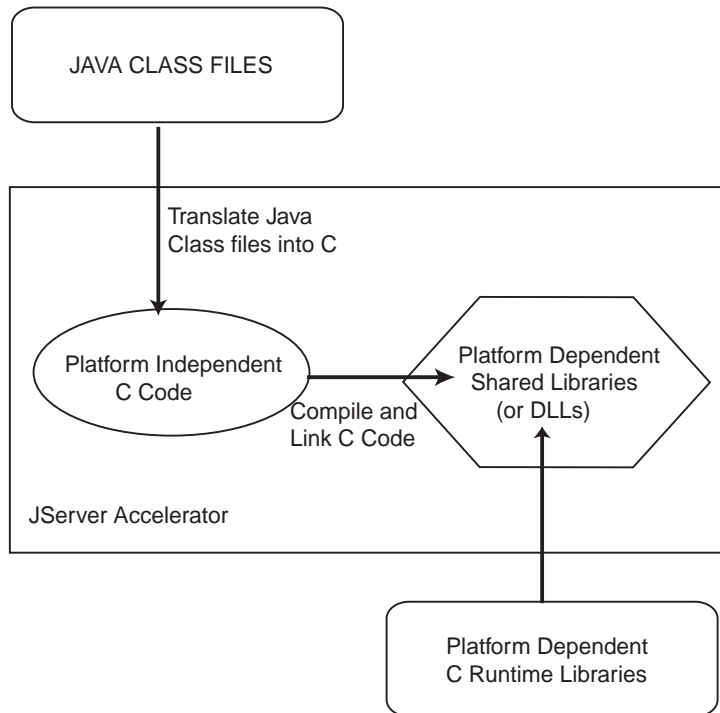
Native Compiler	Description
Just-In-Time	Provides the JVM the ability to translate the Java instructions just before needed by the JDK. The benefits depends on how accurately the native compiler anticipates code branches and the next instruction. If incorrect, no performance gain is realized.

Native Compiler	Description
Ahead-Of-Time	The JServer Accelerator natively compiles all Java code within a JAR file into native shared libraries, which are organized by Java package, before execution time. At runtime, JServer Accelerator checks if a Java package has been natively compiled; and if so, uses the machine code library instead of interpreting the deployed Java code.

This static compilation approach provides a large, consistent performance gain, regardless of the number of users or the code paths they traverse on the server. After compilation, the tool loads the statically compiled libraries into JServer, which are then shared between users, processes, and sessions.

JServer Accelerator Overview

Most Ahead-Of-Time native compilers compile directly into a platform-dependent language. For portability requirements, this was not feasible. As shown in [Figure 6-1](#), the JServer Accelerator translates the Java classes into a version of C that is platform-independent. This C code is compiled and linked to supply the final platform-dependent, natively compiled shared libraries or DLLs.

Figure 6–1 Native Compilation using JServer Accelerator

Given a JAR file, the JServer Accelerator performs the following:

1. The classes, loaded in the database, are verified.
2. The Java bytecodes for these classes are retrieved from the database and stored in a project directory where the JServer Accelerator was invoked.
3. The Java bytecodes are translated to C.
4. The C code is compiled and linked with the C compiler for your platform.

JServer Accelerator translates, compiles, and links the retrieved classes on the client. For this reason, you must natively compile on the intended platform environment that this application will be deployed to. The result is a single deployment JAR file for all classes within the project.

5. The resulting shared library is loaded into the `$ORACLE_HOME/javavm/admin` directory.

Note: The JServer Accelerator natively compiled libraries can only be used within JServer. Also, these libraries can only be used within the same version of JServer that it was produced in. If you want your application to be natively compiled on subsequent releases, you must recompile these classes. That is, native recompilation of existing libraries will not be performed automatically by any upgrade process.

JServer Core Java Class Libraries

All core Java class libraries and Oracle-provided Java code within JServer is natively compiled for greater execution speed. Java classes exist as shared libraries in `$ORACLE_HOME/javavm/admin`, where each shared library corresponds to a Java package. For example, `orajox8java_lang.so` on Solaris and `orajox8java_lang.dll` on Windows NT hold `java.lang` classes. Specifics of packaging and naming can vary by platform. The Aurora JVM uses natively compiled Java files internally and opens them, as necessary, at runtime.

Natively Compiling Java Application Class Libraries

The JServer Accelerator can be used by Java application products that need a performance increase and are deployed on JServer. The JServer Accelerator command-line tool, `ncomp`, natively compiles your code and loads it in JServer. However, in order to use `ncomp`, you must first provide some initial setup.

Installation Requirements

You must install the following before invoking JServer Accelerator:

1. Install a C compiler for the intended platform on the machine you are running `ncomp`.
2. Verify that the correct compiler and linker commands are referenced within the `System*.properties` file located in the `$ORACLE_HOME/javavm/jahome` directory. Since the compiler and linker information is platform-specific, the configuration for these items is detailed in the README for your platform.
3. Grant the user that executes `ncomp` the following role and security permissions:

Note: DBA role contains both the `JAVA_DEPLOY` role and the `FilePermission` for all files under `$ORACLE_HOME`.

- a. `JAVA_DEPLOY`: The user must be assigned to the `JAVA_DEPLOY` role in order to be able to deploy the shared libraries on the server, which both the `ncomp` and `deploync` utilities perform. For example, the role is assigned to DAVE, as follows:

```
SQL> GRANT JAVA_DEPLOY TO DAVE;
```

- b. `FilePermission`: JServer Accelerator stores the shared libraries with the natively compiled code on the server. In order for JServer Accelerator to store these libraries, the user must be granted `FilePermission` for read and write access to directories and files under `$ORACLE_HOME` on the server. One method for granting `FilePermission` for all desired directories is to grant the user the `JAVASYSPRIV` role, as follows:

```
SQL> GRANT JAVASYSPRIV TO DAVE;
```

See the Security chapter in the *Oracle8i Java Developer's Guide* for more information `JAVASYSPRIV` and granting `FilePermission`.

Executing JServer Accelerator

The following sections show how to do basic native compilation using JServer Accelerator. All `ncomp` options are fully described in the *Oracle8i Java Tools Reference*.

Note: Before you natively compile your Java server code, you must have already loaded and tested it within JServer. Native compilation of untested code is not recommended.

Keep in mind that debuggers, such as the debugger provided with JDeveloper, are useful only with interpreted Java code. You cannot debug a natively compiled library.

All of the Java classes contained within a JAR file must already be loaded within the database. Execute the `ncomp` tool to instruct JServer Accelerator to natively compile

all of these classes. The following natively compiles all classes within the `pubProject.JAR` file:

```
ncomp -user scott/tiger pubProject.JAR
```

Note: Because native compilation must compile and link all of your Java classes, this process may execute over the span of a few hours. The time involved in natively compiling your code depends on the number of classes to compile and the type of hardware on your machine.

If you change any of the classes within this JAR file, JServer Accelerator recompiles the shared library for the package that contains the changed classes. It will not recompile all shared libraries. However, if you want all classes within a JAR file to be recompiled—regardless of whether they were previously natively compiled—you execute `ncomp` with the `-force` option, as follows:

```
ncomp -user scott/tiger -force pubProject.JAR
```

For more options, see the *Oracle8i Java Tools Reference*.

Java Memory Usage

The typical and custom database installation process furnishes a database that has been configured for reasonable Java usage during development. However, runtime use of Java should be determined by the usage of system resources for a given deployed application. Resources you use during development can vary widely, depending on your activity. The following sections describe how you can configure memory depending on your performance needs, how to tell how much SGA memory you are using, and what errors denote a Java memory issue:

- [Configuring Memory Initialization Parameters](#)
- [Java Pool Memory](#)
- [Displaying Used Amounts of Java Pool Memory](#)
- [Correcting Out of Memory Errors](#)

Configuring Memory Initialization Parameters

You can modify the following database initialization parameters to tune your memory usage to reflect more accurately your application needs:

- **SHARED_POOL_SIZE**—Shared pool memory is used by the class loader within the JVM. The class loader uses an average of about 8 KB for each loaded class. Shared pool memory is used when loading and resolving classes into the database. It is also used when compiling source in the database or when using Java resource objects in the database.

The memory specified in `SHARED_POOL_SIZE` is consumed transiently when you use `loadjava`. The database initialization process (executing `initjvm.sql` against a clean database, as opposed to the installed seed database) requires `SHARED_POOL_SIZE` to be set to 50 MB as it loads the Java binaries for approximately 8,000 classes and resolves them. The `SHARED_POOL_SIZE` resource is also consumed when you create call specifications and as the system tracks dynamically loaded Java classes at runtime.

- **JAVA_POOL_SIZE**—Aurora's memory manager allocates all other Java state during runtime execution from the amount of memory allocated using `java_pool_size`. This memory includes the shared in-memory representation of Java method and class definitions, as well as the Java objects migrated to session space at end-of-call. In the first case, you will be sharing the memory cost with all Java users. In the second case, under MTS, you must adjust `JAVA_POOL_SIZE` allocation based on the actual amount of state held in static variables for each session. See ["Java Pool Memory"](#) on page 6-9 for more information on `JAVA_POOL_SIZE`.
- **JAVA_SOFT_SESSIONSPACE_LIMIT**—This parameter allows you to specify a soft limit on Java memory usage in a session, which will warn you if you must increase your Java memory limits. The memory that you request in this parameter is allocated when the session is started.

When a user's session-duration Java state exceeds this size, Aurora generates a warning that is written into the trace files. The default is 1 MB. You should understand the memory requirements of your deployed classes, especially as they relate to usage of session space.

- **JAVA_MAX_SESSIONSPACE_SIZE**—If a user-invokable Java program executing in the server can be used in a way that is not self-limiting in its memory usage, this setting may be useful to place a hard limit on the amount of session space made available to it. The default is 4 GB. This limit is purposely set extremely high to be normally invisible.

When a user's session-duration Java state attempts to exceeds this size, your application can receive an out-of-memory failure.

JServer's unique memory management facilities and sharing of read-only artifacts (such as bytecodes) enables HelloWorld to execute with a per-session incremental memory requirement of only 35 KB. More stateful server applications, such as the Aurora/ORB that CORBA and EJB applications use, have a per-session incremental memory requirement of approximately 200 KB. Such applications must retain a significant amount of state in static variables across multiple calls. Refer to the discussion in the section, "[End-of-Call Migration](#)" on page 6-12, for more information on understanding and controlling migration of static variables at end-of-call.

Java Pool Memory

Java pool memory is used in server memory for all session-specific Java code and data within the JVM. Java pool memory is used in different ways, depending on what mode the Oracle8i server is running in.

Java pool memory used within a dedicated server

The shared part of each Java class used per session. This includes readonly memory, such as code vectors, and methods. In total, this can average about 4-8 KB for each class.

None of the per-session Java state of each session. For a dedicated server, this is stored in UGA within the PGA—not within the SGA.

Under dedicated servers, which is probably the case for applications using only Java Stored Procedures, the total required Java pool memory is not much more than 10 MB.

Java pool memory used within a Multi-Threaded Server (MTS)

The shared part of each Java class that is used per session. This includes readonly memory, such as vectors, and methods. In total, this can average about 4-8 KB for each class.

Some of the UGA used for per-session state of each session is allocated from the Java pool memory within the SGA. Since Java pool memory size is fixed, you must estimate the total requirement for your applications and multiply by the number of concurrent sessions they want to create a total amount of necessary Java pool memory. Each UGA grows and shrinks as necessary; however, all UGAs combined must be able to fit within the entire fixed Java pool space.

Under MTS servers, which is the case for applications using CORBA or EJB, this figure could be very large. Java-intensive, multi-user benchmarks could require more than 1 GB. Current size limitations are unknown; however, it is platform dependent.

Note: If you are compiling code on the server, rather than compiling on the client and loading to the server, you might need a bigger `JAVA_POOL_SIZE` than the default 20 MB. EJB deployment uses the Java compiler on the server; therefore, it also requires a larger `JAVA_POOL_SIZE`.

Displaying Used Amounts of Java Pool Memory

You can find out how much of Java pool memory is being used by viewing the `V$SGASTAT` table. Its rows include pool, name, and bytes. Specifically, the last two rows show the amount of Java pool memory used and how much is free. The total of these two items equals what you configured in the database initialization file.

```
SVRMGR> select * from v$sgastat;
POOL          NAME                                BYTES
-----
                fixed_sga                                69424
                db_block_buffers                        2048000
                log_buffer                              524288
shared pool free memory                        22887532
shared pool miscellaneous                      559420
shared pool character set object              64080
shared pool State objects                     98504
shared pool message pool freequeue           231152
shared pool PL/SQL DIANA                      2275264
shared pool db_files                          72496
shared pool session heap                     59492
shared pool joxlod: init P                    7108
shared pool PLS non-lib hp                   2096
shared pool joxlod: in ehe                   4367524
shared pool VIRTUAL CIRCUITS                 162576
shared pool joxlod: in phe                   2726452
shared pool long op statistics array          44000
shared pool table definiti                   160
shared pool KGK heap                         4372
shared pool table columns                    148336
shared pool db_block_hash_buckets            48792
```

shared pool dictionary cache	1948756
shared pool fixed allocation callback	320
shared pool SYSTEM PARAMETERS	63392
shared pool joxlod: init s	7020
shared pool KQLS heap	1570992
shared pool library cache	6201988
shared pool trigger inform	32876
shared pool sql area	7015432
shared pool sessions	211200
shared pool KGFF heap	1320
shared pool joxs heap init	4248
shared pool PL/SQL MPCODE	405388
shared pool event statistics per sess	339200
shared pool db_block_buffers	136000
java pool free memory	30261248
java pool memory in use	19742720

37 rows selected.

Correcting Out of Memory Errors

- [Running out of memory while compiling](#)
- [Running out of memory while loading](#)

Running out of memory while compiling

If you run out of memory while compiling (within `loadjava` or `deployejb`), you should see an error:

```
A SQL exception occurred while compiling: : ORA-04031: unable to allocate bytes
of shared memory ("shared pool","unknown object","joxlod: init h", "JOX: ioc_
allocate_pal")
```

The cure is to shut down your database and to reset `JAVA_POOL_SIZE` to a larger value. The mention of "shared pool" in the error message is a misleading reference to running out of memory in the "Shared Global Area". It does not mean you should increase your `SHARED_POOL_SIZE`. Instead, you must increase your `JAVA_POOL_SIZE`, restart your server, and try again.

Running out of memory while loading

If you run out of memory while loading classes, it can fail silently, leaving invalid classes in the database. Later, if you try to invoke or resolve any invalid classes, you will see `ClassNotFoundException` or `NoClassDefFoundException`

exceptions being thrown at runtime. You would get the same exceptions if you were to load corrupted class files. You should perform the following:

- Verify that the class was actually included in the set you are loading to the server. Many people have accidentally forgotten to load just one class out of hundreds and spend considerable time chasing this down.
- Use the `loadjava -force` option to force the new class being loaded to replace the class already resident in the server.
- Use the `loadjava -resolve` option to attempt resolution of a class during the load process. This allows you to catch missing classes at load time, not run time.
- Double check the status of a newly loaded class by connecting to the database in the schema containing the class, and execute the following:

```
select * from user_objects where object_name = dbms_java.shortname('');
```

The STATUS field should be "VALID". If `loadjava` complains about memory problems or failures such as "connection lost", increase `SHARED_POOL_SIZE` and `JAVA_POOL_SIZE`, and try again.

End-of-Call Migration

Aurora preserves the state of your Java program between calls by migrating all objects reachable from static variables into session space at the end of the call. Session space exists within the client's session to store static variables and objects that exist between calls. Aurora performs this migration operation at the end of every call, without any intervention by you.

This migration operation is a memory and performance consideration; thus, you should be aware of what you designate to exist between calls and keep the static variables and objects to a minimum. If you store objects in static variables needlessly, you impose an unnecessary burden on the memory manager to perform the migration and consume per-session resources. By limiting your static variables to only what is necessary, you help the memory manager and improve your server's performance.

To maximize the number of users who can execute your Java program at the same time, it is important to minimize the footprint of a session. In particular, to achieve maximum scalability, an inactive session should take up as little memory space as possible. A simple technique to minimize footprint is to release large data structures at the end of every call. You can lazily recreate many data structures when you need them again in another call. For this reason, the Aurora JVM has a mechanism for

calling a specified Java method when a session is about to become inactive, such as at end-of-call time.

This mechanism is the `EndOfCallRegistry` notification. It enables you to clear static variables at the end of the call and reinitialize the variables using a lazy initialization technique when the next call comes in. You should execute this only if you are concerned about the amount of storage you require the memory manager to store in between calls. It becomes a concern only for more complex stateful server applications you implement in Java.

The decision of whether to null-out data structures at end-of-call and then recreate them for each new call is a typical time and space trade-off. There is some extra time spent in recreating the structure, but you can save significant space by not holding on to the structure between calls. In addition, there is a time consideration because objects—especially large objects—are more expensive to access after they have been migrated to session space. The penalty results from the differences in representation of session, as opposed to call-space based objects.

Examples of data structures that are candidates for this type of optimization include:

- Buffers or caches.
- Static fields, such as Arrays, that once initialized can remain unchanged during the course of the program.
- Any dynamically built data structure that could have a space efficient representation between calls and a more speed efficient representation for the duration of a call. This can be tricky and complicate your code, making it hard to maintain, so you should consider doing this only after demonstrating that the space saved is worth the effort.

Oracle-Specific Support for End-of-Call Optimization

You can register the static variables that you want cleared at the end of the call when the buffer, field, or data structure is created. Within the Oracle-specified `oracle.aurora.memoryManager.EndOfCallRegistry` class, the `registerCallback` method takes in an object that implements a `Callback` object. The `registerCallback` object stores this object until the end of the call. When end-of-call occurs, Aurora invokes the `act` method within all registered `Callback` objects. The `act` method within the `Callback` object is implemented to clear the user-defined buffer, field, or data structure. Once cleared, the `Callback` is removed from the registry.

Note: If the end of the call is also the end of the session, callbacks are not invoked, because the session space will be cleared anyway.

The way you use the `EndOfCallRegistry` depends on whether you are dealing with objects held in static fields or instance fields.

- **Static fields**—You use `EndOfCallRegistry` to clear state associated with an entire class. In this case, the `Callback` object should be held in a private static field. Any code that requires access to the cached data dropped between calls must invoke a method that lazily creates—or recreates—the cached data. The example below does the following:
 1. Creates a `Callback` object within a static field, `thunk`.
 2. Registers this `Callback` object for end-of-call migration.
 3. Implements the `Callback.act` method to free up all static variables, including the `Callback` object itself.
 4. Provides a method, `createCachedField`, for lazily recreating the cache.

When the user creates the cache, the `Callback` object is automatically registered within the `getCachedField` method. At end-of-call, Aurora invokes the registered `Callback.act` method, which frees the static memory.

```
import oracle.aurora.memoryManager.Callback;
import oracle.aurora.memoryManager.EndOfCallRegistry;

class Example {
    static Object cachedField = null;
    private static Callback thunk = null;

    static void clearCachedField() {
        // clear out both the cached field, and the thunk so they don't
        // take up session space between calls
        cachedField = null;
        thunk = null;
    }

    private static Object getCachedField() {
        if (cachedField == null) {
            // save thunk in static field so it doesn't get reclaimed
            // by garbage collector
            thunk = new Callback () {

```



```

        public void act(Object obj) {
            Example.clearCachedField();
        }
    };

    // register thunk to clear cachedField at end-of-call.
    EndOfCallRegistry.registerCallback(thunk);
    // finally, set cached field
    cachedField = createCachedField();
}
return cachedField;
}

private static Object createCachedField() {
    ....
}
}

```

- Instance fields—Use `EndOfCallRegistry` to clear state in data structures held in instance fields. For example, when a state is associated with each instance of a class, each instance has a field that holds the cached state for the instance and fills in the cached field as necessary. You can access the cached field with a method that ensures the state is cached.
 1. Implements the instance as a `Callback` object.
 2. Implements the `Callback.act` method to free up the instance's fields.
 3. When the user requests a cache, the `Callback` object registers itself for end-of-call migration.
 4. Provides a method, `createCachedField`, for lazily recreating the cache.

When the user creates the cache, the `Callback` object is automatically registered within the `getCachedField` method. At end-of-call, Aurora invokes the registered `Callback.act` method, which frees the cache.

This approach ensures that the lifetime of the `Callback` object is identical to the lifetime of the instance, because they are the same object.

```

import oracle.aurora.memoryManager.Callback;
import oracle.aurora.memoryManager.EndOfCallRegistry;

class Example2 implements Callback {
    private Object cachedField = null;

    public void act(Object obj) {

```

```
        // clear cached field
        cachedField = null;
        obj = null;
    }

    // our accessor method
    private static Object getCachedField() {
        if (cachedField == null) {
            // if cachedField is not filled in then we need to
            // register self, and fill it in.
            EndOfCallRegistry.registerCallback(self);
            cachedField = createCachedField();
        }
        return cachedField;
    }

    private Object createCachedField() {
        ....
    }
}
```

A weak table holds the registry of end-of-call callbacks. If either the `Callback` object or value are not reachable (see JLS section 12.6) from the Java program, they will both be dropped from the table. The use of a weak table to hold callbacks also means that registering a callback will not prevent the garbage collector from reclaiming that object. Therefore, you must hold on to the callback yourself if you need it—you cannot rely on the table holding it back.

You can find other ways in which end-of-call notification will be useful to your applications. The following sections give the details for methods within the `EndOfCallRegistry` class and the `Callback` interface:

EndOfCallRegistry.registerCallback method

The `registerCallback` method installs a `Callback` object within a registry. At the end of the call, Aurora invokes the `act` methods of all registered `Callback` objects.

You can register your `Callback` object by itself or with a value object. If you need additional information stored within an object to be passed into `act`, you can register this object within the `value` parameter.

```
public static void registerCallback(Callback thunk, Object value);
public static void registerCallback(Callback thunk);
```

Parameter	Description
thunk	The Callback object to be invoked at end-of-call migration.
value	If you need additional information stored within an object to be passed into <code>act</code> , you can register this object within the <code>value</code> parameter. In some cases, the <code>value</code> parameter is necessary to hold state the callback needs. However, most users do not need to specify a <code>value</code> .

EndOfCallRegistry.runCallbacks method

```
static void runCallbacks()
```

The JVM calls this method at end-of-call and calls `act` for every `Callback` object registered using `registerCallback`. You should never call this method in your code. It is called at end-of-call, before object migration and before the last finalization step.

Callback Interface

```
Interface oracle.aurora.memoryManager.Callback
```

Any object you want to register using

`EndOfCallRegistry.registerCallback` implements the `Callback` interface. This interface can be useful in your application, where you require notification at end-of-call.

Callback.act method

```
public void act(Object value)
```

You can implement any activity that you require to occur at the end of the call. Normally, this method will contain procedures for clearing any memory that would be saved to session space.

Memory Profiling Utility

The purpose of the Memory Profiling Utility (MemStat) is to trace, profile, and report on the allocated memory that is accessible through static variables in your Oracle8i Java program. You can then use the information in this report to locate and eliminate unnecessary static data in your Java classes, thereby reducing the static footprint of your Java program and improving the performance of repeated Java calls into the database.

The Oracle8i JVM uses three kinds of memory:

- *call memory*, which exists for the duration of a (possibly recursive) call into the database
- *session memory*, which exists for the duration of the session connected to the database
- *permanent or global memory*, which persists as long as the database instance is running

Java language semantics specify that static variables persist across calls. At the end of each call, the Oracle8i JVM copies the *call* memory that is accessible through the static variables in each class into *session* memory so that it can be saved and restored on subsequent calls to methods in those Java classes. If there is a lot of static data or a complex graph of interconnected objects, then there is considerable overhead during the end-of-call processing while the JVM allocates session memory and copies the static data to it.

A typical technique for tuning object-oriented programs for faster performance is to eliminate the allocation of unnecessary objects from your program. For example, you can create a static instance of a commonly used object and reuse it rather than creating a new one every time you need it. However, the interactions among the different database memories complicate such techniques, and can require analysis of the speed trade-off for allocating dynamic objects versus the space trade-off for the end-of-call copying of static objects. If a static object is large, or if there are many such objects, or if there are many calls, then the speed advantage gained by caching the object may be lost, due to the traversal of the object graph during end-of-call processing.

How MemStat Works

Depending on how you invoke it, MemStat will analyze either a single class or all classes that are loaded into the current session. For each class, MemStat enumerates the static variables of the class. These variables are known as the *roots*. Depending on the structure of each variable, MemStat performs three different analyses:

- If the variable is a primitive object, such as an integer, MemStat records its class and size.
- If it is a non-primitive object (for example, one that refers to another object), MemStat follows the reference and recursively enumerates all objects to which it refers.
- If it is an array, MemStat enumerates all elements of the array.

This process is repeated recursively until all objects reachable from all static variables have been recorded. Because it is possible for large object graphs to contain cycles, MemStat also records any circular references it encounters during the analysis.

Using MemStat

The purpose of MemStat is to analyze and report on the object graph accessible from the static variables in your program. You can invoke the analysis directly from any point in your program, and you can also register it to run at the end of a call.

Because there is no *standard* output mechanism for database calls, MemStat produces its report in the form of HTML files, in a directory you specify. When the report is finished, you can view these files with any HTML-capable Web browser.

MemStat is implemented in three static methods on the class `oracle.aurora.memstat.MemStat`.

You can call it in three different ways:

- to report immediately on the static variables in a single class
- to report immediately on the static variables of all classes used in the current session
- to report at end-of-call on the static variables of all classes used in the current session

The method call for reporting on a single class is:

```
MemStat.writeDump (Class MyClass, String outputPath, String filePrefix);
```

The method call for reporting on all loaded classes is:

```
MemStat.writeDump (String outputPath, String filePrefix);
```

The method call for reporting on all loaded classes at the end-of-call is:

```
MemStat.writeDumpAtEOC (String outputPath, String filePrefix);
```

The `outputPath` parameter represents the directory in which the MemStat reports are generated. The `outputPath` string must be in a file name format that is suitable to the platform on which the report is generated. For example, `/home/base/memstat` is suitable for a Solaris platform; the Windows format might be `c:\\base\\memstat`. Note that Java requires doubling of the backslashes inside a string, but not the forward slashes.

The `filePrefix` is the base file name for the HTML files that are generated in the `outputPath` directory. Because MemStat reports can be voluminous, and many Web browsers have limitations on the size of the files they can browse, MemStat breaks long reports into separate files. The `filePrefix` is the basis for all file names in a given report and is augmented by an incremental numeric suffix. If, for example, the `test` report produces three files, the main report file will be named `test.htm`, and additional report files will be named `test1.htm` and `test2.htm`.

If you call MemStat more than once in a given call, be careful to use different base names or different output directories, lest the subsequent reports overwrite the previous ones. For example, if you call MemStat before and after you perform some memory-consuming operation, naming the first report `before` and the second report `after` will prevent name collisions, while still writing the report files into the same directory. Using multiple directories is more complicated: you must remember to grant separate FilePermissions (see below) for each directory in which you want to write.

Here are some sample MemStat calls:

```
MemStat.writeDump (MyClass.class, "c:\\base\\memstat", "myclass");
MemStat.writeDump ("/home/base/memstat", "test");
MemStat.writeDumpAtEOC ("/home/base/memstat", "eoc");
```

MemStat Permissions

MemStat requires certain permissions to be granted to the user or role under which it runs. Because MemStat runs in the Oracle server process, these permissions grant access to the resources that MemStat requires:

- Access to the private variables of the objects in the graph. If this permission is not granted, MemStat will still run, but it will not trace any objects pointed to by private variables.
- Access to the Java run-time system for determining which classes are loaded in the current session. If this permission is not granted, MemStat will find zero classes loaded and will generate an empty report.
- Access to the server file system where the MemStat HTML reports are generated. If this permission is not granted, MemStat will raise an exception when it tries to create the report files.

The following SQL statements grant these permissions to user JIM:

```
call dbms_java.grant_permission ('JIM',
    'SYS:java.lang.reflect.ReflectPermission', 'suppressAccessChecks', null);
```

```
call dbms_java.grant_permission ('JIM',
                                'SYS:oracle.aurora.security.JServerPermission', 'JRIExtensions', null);

call dbms_java.grant_permission ('JIM', 'SYS:java.io.FilePermission',
                                '/home/base/memstat', 'read,write');           // Solaris

call dbms_java.grant_permission ('JIM', 'SYS:java.io.FilePermission',
                                'c:\base\memstat', 'read,write');           // Windows
```

If the Oracle Server is running on a Windows platform, the output file path named in the MemStat call is subtly different from the path in the SQL `grant_permission` call. In Java strings, you must use double backslashes; in SQL you need only one backslash.

The MemStat Report Format

This section describes the format of the MemStat report; see the example below. You can browse the MemStat output report with any HTML-capable Web browser. To do this, point the browser at the base file name that is specified, for example, `c:\base\memstat\test.htm`.

The report begins with a summary of the memory usage at the time MemStat is invoked. This summary should give you an overall idea of the amount of memory used by the Java call that you are analyzing.

Following the summary is a list of the unique classes that are traversed during the MemStat analysis. For every object found during the memory analysis, MemStat has recorded its class and its size in each of call, session, and permanent memory. The largest objects are sorted first, because eliminating these will yield the largest performance improvement. The list is actually sorted by the largest of these three sizes, calculated as *max (call, max (session, permanent))*. For each class, this table also shows how many bytes are occupied by objects of that class, how many objects there are, their minimum, maximum and average size, and for arrays, the standard deviation of the sizes.

Following the class summary is one or more tables describing each *root* object. The title of the object describes the package and class of the object. Each row of the table describes:

- a field of the object
- a description of the object in that field
- the total size of the object in each of the three memory spaces

Following the root objects are the objects pointed to by the roots; the objects are separated by a dividing rule. One, two, or three tables describe each object:

- the object itself
- any other objects that the object refers to
- any objects that refer to this object

The title for each object describes the memory in which the object resides: Call, Session, or Permanent. Each object is described by:

- a unique identifier
- the output of applying the `toString()` method to the object
- the object's direct size: the size of the object header and the fields of the object
- the object's total size: the sum of the sizes of all the objects to which it refers

An object that refers to another object is linked by an HTML link to the tables representing the object to which it refers. You can navigate the object graph using these links as you would navigate hyperlinks in a text document.

Sample Output

MemStat Results

2000-06-01 17:07:05.645

Run-Time Values			
Session Size	143360		
NewSpace Size	262144		
NewSpace Enabled	true		
Intern Table Size	261814		

Total Memory Allocation			
	Call	Session	Permanent
Objects	726	926	3217
Total Size	54861	39348	127418
Minimum	12	12	12

Total Memory Allocation			
Maximum	16396	2060	8076
Average	75.6	42.5	39.6
Std Deviation	679.2	93.7	233.7

Allocated Objects by Class (Call, Session, Permanent)																			
Class	Bytes			Objects			Minimum		Maximum			Average		Standard Deviation					
char[]	25316	5134	43296	161	104	1177	2	0	0	16384	800	8064	157.2	49.4	36.8	1,283.2	105.0	276.9	
java.lang.String	3816	3240	30528	159	135	1272													
java.util.Hash table\$Entry	4956	10696	12460	177	382	445													
byte[]	8195	2421	2107	2	34	57	3	0	0	8192	2048	1024	4,097.5	71.2	37.0	4,094.5	344.2	143.3	

Objects Accessible From java.util.Properties		
Field	Reference	Total Size
keyValueSeparators	946: java.lang.String	0 0 50
strictKeyValueSeparators	948: java.lang.String	0 0 40
specialSaveChars	950: java.lang.String	0 0 54
whiteSpaceChars	952: java.lang.String	0 0 46
hexDigit	954: char[16]	0 0 44

Glossary

API

Application Programming Interface. As applied to Java, a well-defined set of classes and methods that furnish a specific set of functionality to the Java programmer. JDBC and SQLJ are APIs for accessing SQL data.

Bytecodes

The set of single-byte, machine-independent instructions to which Java source code is compiled using the Java compiler.

Call Memory

The memory that the memory manager uses to allocate new objects.

CLASSPATH

The environment variable (or command line argument) the JDK or JRE uses to specify the set of directory tree roots in which Java source, classes, and resources are located.

Context switch

In a uniprocessor system, the current thread is interrupted by a higher priority thread or by some external event, and the system switches to a different thread. The choice of which thread to dispatch is usually made on a priority basis or based on how long a thread has been waiting.

Cooperative Multitasking

The programmer places calls to the `Thread.yield()` method in locations in the code where it is appropriate to suspend execution so that other threads can run.

This is quite error-prone because it is often difficult to assess the concurrent behavior of a program as it is being written.

CORBA

Common Object Request Broker Architecture. Specified by the Object Management Group (OMG), CORBA provides a language-independent architecture for distributing object-oriented programming logic between logical and physical tiers in a network, connected through ORBs.

Core Class Libraries

Generally, the Java packages delivered with Sun Microsystem's JDK, java.*. We also use this term to denote some sun.* packages.

Deadlock

The conflict state where two or more synchronized Java objects depend on locking each other but cannot because they themselves are locked by the dependent object. For example, object A tries to lock object B while object B is trying to lock object A. This situation is difficult to debug because a preemptive Java virtual machine can neither detect nor prevent deadlock. Without deadlock detection, a deadlocked program simply hangs.

Dispatch

The system saves the state of the currently executing thread, restores the state of the thread to be executed, and branches to the stored program counter for the new thread, effectively continuing the new thread as if it had not been interrupted.

Driver

As used with JDBC, a layer of code that determines the low-level libraries employed to access SQL data and/or communicate across a network. The three JDBC drivers in JServer are: thin, OCI, and server.

EJB

Enterprise JavaBeans. JServer provides an implementation of the Enterprise JavaBeans 1.0 Specification. JServer supports only Session Beans; Entity Beans are an optional part of the EJB 1.0 Specification, and JServer does not support them.

End-of-Call

Within your session, you may invoke Java many times. Each time you perform this, end-of-call occurs at the point at which Java code execution completes. The memory manager migrates static variables to session space at end-of-call.

Garbage Collection

The popular name for the automatic storage reclamation facility provided by the Java virtual machine.

IDE

Integrated Development Environment. A Java IDE runs on a client workstation, providing a graphical user interface for access to the Java class library and development tools.

Interface Definition Language (IDL)

The platform-independent language CORBA specifies for defining the interface to a CORBA component. You use a tool like `idl2java` to convert IDL to Java code.

Java Schema Object

The term JServer uses to denote either Java source, binary, or resources when stored in the Oracle8i database. These three Java schema objects correspond to files under the JDK—`.java`, `.class`, or other files (such as `.properties` files) used in the JDK CLASSPATH.

JCK

Java Compatibility Kit. The set of Java classes that test a Java virtual machine and Java compiler's compliance with the Java standard. JCK releases correspond to Sun Microsystems's JDK releases, although in the case of JServer, only the Java classes and not the virtual machine, are identical to Sun Microsystems's JDK.

JDBC

Java Database Connectivity. The standard Java classes that provide vendor-independent access to databases.

JDBC Driver

The vendor-specific layer of JDBC that provides access to a particular database. Oracle provides three JDBC drivers—`thin`, `OCI`, and `server`.

JDK

Java Development Kit. The Java virtual machine, together with the set of Java classes and tools Sun Microsystems furnishes to support Java application and applet development. The JDK includes a Java compiler; the JRE does not.

JRE

Java Runtime Environment. The set of Java classes supporting a Java application or applet at runtime. The JRE classes are a subset of the JDK classes.

JServer

Oracle's scalable Java server platform, composed of the Aurora Java virtual machine running within the Oracle8i database server, the Java runtime environment and Oracle extensions, including the Aurora/ORB and Enterprise JavaBeans implementation.

Lazy Initialization

A technique for initializing data, typically used in accessor methods. The technique checks to see if a field has been initialized (is non-null) before returning the initialized object to it. The overhead associated with the check is often small, especially in comparison to initializing a data structure that may never be accessed. You can employ this technique in conjunction with end-of-call processing to minimize session space overhead.

Object Graph

An object is said to reference the objects held in its fields. This collection of objects forms an object graph. The memory manager actually migrates the object graphs held in static variables; that is, it migrates not only the objects held in static fields, but the objects that those objects reference, and so on.

ORB

Object Request Broker. An ORB is a program that executes on the server, receiving encoded messages from clients for execution by server-side objects and returning objects to the client. ORBs typically support different services that clients can use, such as a name service. The JServer ORB is known as the Aurora/ORB.

Preemptive Multitasking

The operating system preempts, or takes control away from a thread, under certain conditions, such as when another thread of higher priority is ready to run, or when an external interrupt occurs, or when the current thread waits on an I/O operation, such as a socket accept or a file read. Some Java virtual machines implement a type of round-robin preemption by preempting the current thread on certain virtual machine instructions, such as backward branches, method calls, or other changes in control flow. For a Java virtual machine that maps Java threads to actual operating system threads, the preemption takes place in the operating system kernel, outside

the control of the virtual machine. Although this yields decent parallelism, it complicates garbage collection and other virtual machine activities.

Process

An address space and one or more threads.

Session Memory

The memory that the memory manager uses to hold objects that survive past the end-of-call—those objects reachable from Java static variables within your session.

SQLJ

Embedded SQL in Java. The standard that defines how SQL statements can be embedded in Java programs to access SQL data. A translator transforms the SQLJ programs to standard JDBC programs.

Strong Typing

In Java, the requirement that the class of each field and variable, and the return type of each method be explicitly declared.

Symmetric Multiprocessing (SMP)

The hardware has multiple processors, and the operating system maps threads to different processors depending on their load and availability. This assumes that the Java virtual machine maps OS threads to Java threads. This mechanism provides true concurrency among the threads but can lead to subtle programming errors and deadlock conflicts on synchronized objects.

System

Often used in discussion as the combination of the hardware, the operating system and the Java virtual machine.

Thread

An execution context consisting of a set of registers, a program counter, and a stack.

Virtual Machine

A program that emulates the functionality of a traditional processor. A Java virtual machine must conform to the requirements of the Java Virtual Machine Specification. The JServer virtual machine is known as the Aurora virtual machine.

Symbols

#sql, 3-13, 3-14

A

act method, 6-13

application

- compiling, 2-15
- development, 2-3
- distributed, 1-21
- executing in a session, 2-3
- execution control, 2-6
- execution rights, 2-27
- invoking, 3-3, 3-24
- threading, 2-44

attributes

- definition, 1-3
- types of, 1-4

Aurora

- definition, 2-2

authentication, 5-2

B

BasicPermission, 5-14

bean, 1-21

bytecode

- defined, 1-8
- verification, 2-22

C

call

definition, 2-2

managing resources across calls, 2-41

static fields, 2-5

call specification, 3-5

Callback class

act method, 6-13

class

- attributes, 1-3, 1-5
- definition, 1-2
- dynamic loading, 1-16
- execution, 2-2
- hierarchy, 1-5
- inheritance, 1-5, 1-6
- loading, 2-2, 2-6, 2-23
- marking valid, 2-20
- methods, 1-3, 1-5
- name, 2-32
- publish, 2-2
- resolving references, 2-19

Class class

getClassLoader method, 2-35

.class files, 2-14, 2-23, 2-24

Class interface

forName method, 2-33

class schema object, 2-14, 2-20, 2-23, 2-24

classes

- loading, 3-2
- protected, 5-25
- publishing, 2-30, 3-2
- resolving, 3-2

classes111.zip, 2-9

classes12.zip, 2-10

Class.forName class

lookupClass method, 2-36

classForNameAndSchema method, 2-36

ClassNotFoundException, 2-33

CLASSPATH, 2-14, 2-33, 4-9

client

setup, 4-8

code

native compilation, 6-2

CodeSource class, 5-5

equals method, 5-5

implies method, 5-5

COM, 1-22

compiling, 2-15

error messages, 2-16

memory problems, 6-11

options, 2-16

runtime, 2-15

setting options, 3-19

component, 1-21

configuration, 4-1

JServer, 4-6 to 4-8

performance, 6-7

connection

configuration, 4-6

security, 5-2

CORBA

configuring, 4-7, 4-8

CosNaming, 3-7

defined, xii, 1-17, 1-21, 2-5, 3-2, 3-6

documentation, 1-22

example, 3-9

invoking, 3-2

Java 2 support, 2-10

pure CORBA using Java 2, 2-13

security, 5-2

CosNaming, 3-7

D

data confidentiality, 5-2

database

configuration, 4-6

privileges, 5-2

DBA_JAVA_POLICY view, 5-6, 5-18, 5-20

DBMS_JAVA package, 3-21, 4-3

defined, 5-5

delete_permission method, 4-5, 5-18

disable_permission method, 4-5, 5-18

dropjava method, 4-4

enable_permission method, 4-5, 5-18

get_compiler_option method, 4-4

grant_permission method, 4-5, 5-8, 5-10

grant_policy_permission method, 4-5, 5-12, 5-20

loadjava method, 4-4

longname method, 2-29, 2-33, 4-3

manipulating security, 2-10

modifying permissions, 5-19

modifying PolicyTable permissions, 5-10, 5-12

reset_compiler_option method, 4-4

restart_debugging method, 3-22, 4-5

restrict_permission method, 4-5, 5-9, 5-10

revoke_permission method, 4-5, 5-18

set_compiler_option, 3-19

set_compiler_option method, 4-4

set_output method, 3-24, 4-4

setting permissions, 5-6

shortname method, 2-29, 2-33, 4-3

start_debugging method, 3-21, 4-5

stop_debugging method, 4-5

DBMS_OUTPUT package, 4-4

DbmsJava class, see DBMS_JAVA package

DbmsObjectInputStream class, 2-36

DbmsObjectOutputStream class, 2-36

deadlock, 2-45

DeadlockError exception, 2-45

debug

compiler option, 2-17

DebugAgent class, 3-19

debugging, 4-5, 5-25

agent, 3-19, 3-21

connecting a debugger, 3-22

Java applications, 3-18

necessary permissions, 5-25

setting compiler options, 3-19

starting Debug Agent, 3-21

starting proxy, 3-20

using OracleAgent class, 3-22

DebugProxy class, 3-19, 3-20

definer rights, 2-27

delete method, 5-18

delete_permission method, 4-5, 5-18

- deployejb tool, 2-31
- disable method, 5-18
- disable_permission method, 4-5, 5-18
- distributed objects, 1-21, 3-6
- dropjava method, 4-4
- dropjava tool, 2-25

E

EJB

- component, 1-21
- configuring, 4-7, 4-8
- defined, xii, 1-17, 1-21, 2-5, 3-2, 3-6
- documentation, 1-22
- example, 3-8
- invoking, 3-2
- security, 5-2
- state, 1-21

enable method, 5-18

enable_permission method, 4-5, 5-18

encoding

- compiler option, 2-17

end-of-call migration, 6-12

EndOfCallRegistry class, 6-13

- registerCallback method, 6-13

endSession method, 2-5, 2-46

Enterprise Java Beans, see EJB

equals method, 5-5

errors

- compilation, 2-16

exception

- ClassNotFoundException, 2-33
- DeadlockError, 2-45
- IOException, 2-39
- LimboError, 2-45
- ThreadDeathException, 2-46

execution rights, 2-27

exitCall method, 2-46

F

- FilePermission, 5-8, 5-19, 5-21, 5-23
- FilePermission permission, 6-6
- files, 2-38
 - across calls, 2-41

- lifetime, 2-39

finalizers, 2-41

footprint, 1-14, 2-4

forName method, 2-33

G

garbage collection, 1-13, 1-14, 2-5

- managing resources, 2-38
- misuse, 2-40
- purpose, 2-40

General Inter-Orb Protocol, see GIOP

get_compiler_option method, 2-17, 4-4

getCallerClass method, 2-35

getClassLoader method, 2-35

GIOP

- configuring, 4-6, 4-7
- presentation, 4-8

grant method, 5-8

grant_permission method, 4-5, 5-8, 5-10

grant_policy_permission method, 4-5, 5-12, 5-20

granting permission, 5-5

grantPolicyPermission method, 5-13

GUI, 2-31

I

IIOP

- configuring, 4-7
- defined, 3-7
- SSL, 4-8

implies method, 5-5

inheritance, 1-5, 1-6

init method, 2-11

initjvm.sql, 4-2, 4-3, 4-6

installation, 4-1, 4-2

integrity, 5-2

interfaces

- defined, 1-6
- user, 2-31

internet newsgroups, xiii

invoker rights, 2-27

IOException, 2-39

J

Java

- applications, 2-1, 2-15
 - loading, 2-23
 - attributes, 1-3
 - class, 1-2
 - client
 - CLASSPATH, 4-9
 - setup, 4-8
 - compiling, 2-15
 - development environment, 2-14
 - differences from Sun JDK, 2-3
 - distributed applications, xii, 3-2
 - documentation, xii, 1-1, 1-22
 - execution control, 2-6
 - execution rights, 2-27
 - features, 1-11
 - in the database, 1-1, 1-12, 2-1, 2-2
 - interpreter, 2-2
 - introduction, xi
 - invoking, 2-2, 3-3
 - loading classes, 2-6, 3-2
 - checking results, 2-28
 - methods, 1-3
 - natively compiling, 6-2
 - overview, 1-1, 1-2
 - permissions, 4-5
 - polymorphism, 1-6
 - programming models, xii
 - publishing, 2-7
 - resolving classes, 2-19
 - resources, 1-2
 - stored procedures, see Java stored procedures
- Java 2
- migrating from JDK 1.1, 2-7
 - migrating security, 2-10
 - security, 5-2
- Java Compatibility Kit, see JCK
- .java files, 2-14, 2-23, 2-24
 - java interpreter, 2-2, 2-6
- Java language specification, see JLS
- Java Naming and Directory Interface, see JNDI
- Java Native Interface, see JNI
- Java Remote Method Invocation, see RMI

- Java stored procedures, xii, 2-5
 - configuring, 4-7
 - defined, 1-17, 1-18, 3-3
 - documentation, 1-22
 - invoking, 3-2
 - publishing, 2-30
- Java virtual machine, see JVM
- JAVASOPTIONS table, 2-16
- JAVA_ADMIN role
 - assigned permissions, 5-21
 - example, 5-13
 - granting permission, 5-3, 5-5, 5-12, 5-20
- JAVA_DEPLOY role, 6-6
- JAVA_MAX_SESSIONSPACE_SIZE
 - parameter, 6-8
- JAVA_POOL_SIZE parameter
 - default, 4-6
 - defined, 6-8, 6-9
 - errors, 6-11
 - minimum value, 4-3
- JAVA_SOFT_SESSIONSPACE_LIMIT
 - parameter, 6-8
- JAVADEBUGPRIV role, 5-24, 5-25
- JavaServer Pages, 3-2
- java.sql package, 2-8
- JAVASYSPRIV role, 5-3, 5-23, 5-24
- JAVAUERPRIV role, 5-3, 5-23, 5-24
- JCK, 1-11
- jdb debugging tool, 3-18, 3-23
- JDBC
 - 2.0 support, 2-8
 - accessing SQL, 1-18
 - defined, 1-17, 3-2, 3-11
 - documentation, 1-22
 - driver types, 1-19, 3-11
 - example, 3-12
 - interacting with SQL, 3-17
 - security, 5-2
 - web information, xii
- JDeveloper
 - debugging, 3-18
 - development environment, 1-22, 3-16, 4-12
- JDK
 - requirements, 4-8
 - web location, xiii

- JLS
 - specification, 1-11
 - web information, xiii
- JNDI
 - defined, 3-7
 - name lookup, 3-8
- JNI support, 3-10
- JPublisher
 - documentation, 1-22
- JServer
 - configure, 4-1
 - definition, xiii, 2-2
 - install, 4-1, 4-2
- JServer Accelerator
 - for user applications, 6-5
 - installation requirements, 6-5
 - overview, 6-2, 6-3
- JServerPermission, 5-8, 5-19, 5-20, 5-21, 5-22, 5-23, 5-24
 - defined, 5-20
- JVM
 - bytecodes, 1-8
 - defined, 1-2, 1-8
 - garbage collection, 1-13, 1-14
 - multithreading, 1-13
 - responsibilities, 2-4
 - security, 4-5
 - specification, 1-11
 - web information, xiii

L

- LimboError exception, 2-45
- loading, 2-23 to 2-30
 - checking results, 2-25, 2-28
 - class, 1-16, 2-6, 2-15
 - compilation option, 2-15
 - granting execution, 2-27
 - JAR or ZIP files, 2-26
 - necessary privileges and permissions, 2-26
 - reloading classes, 2-27
 - restrictions, 2-25
- loadjava method, 4-4
- loadjava tool, 2-24 to 2-26
 - compiling source, 2-15, 6-12

- example, 3-4
- execution rights, 2-27, 5-3
- loading class, 2-23
- loading ZIP or JAR files, 2-26
- restrictions, 2-25
- using memory, 6-8
- logging, 2-16
- longname method, 2-29, 2-33, 4-3
- lookupClass method, 2-36

M

- main method, 2-6
- memory
 - across calls, 2-40
 - call, 2-5
 - java pool, 6-10
 - leaks, 2-40
 - lifetime, 2-38, 2-39
 - manager, 2-14
 - performance configuration, 6-7
 - report allocation, 6-17
 - running out of, 6-11
 - session, 2-5, 6-13
- Memory Profiling Utility, see MemStat
- MemStat
 - analysis options, 6-18
 - example, 6-20, 6-22
 - reporting, 6-21
 - security permissions, 6-20
 - using, 6-19
 - warning, 6-20
- MemStat class, 6-19
 - writeDump method, 6-19
 - writeDumpAtEOC method, 6-19
- MemStat utility
 - overview, 6-17
- methods, 1-3, 1-5
- multithreading, 1-13

N

- name service, 3-7
 - CosNaming, 3-7
 - JNDI, 3-7

- native compilation, 1-15, 6-2
 - static, 6-3
- native compiler, 6-2
- ncomp
 - security, 6-6
- ncomp tool, 6-5
 - executing, 6-7
- Net8
 - configuring, 4-6
- NetPermission, 5-8, 5-19, 5-21, 5-22
- networking
 - configuration, 4-6

O

- object
 - full to short name conversion, 2-29
 - lifetime, 2-39
 - schema, 2-14
 - serialization, 2-36
 - short name, 2-29
- ObjectInputStream class, 2-36
- ObjectOutputStream class, 2-36
- online
 - compiler option, 2-17
- operating system
 - resources, 2-38
 - across calls, 2-41
 - lifetime, 2-39
 - performance, 6-7
 - permission, 2-39
- OracleAgent class
 - restart method, 3-22
 - start method, 3-22
 - stop method, 3-22
- oracle.jdbc2 package, 2-8, 2-10
- OracleRuntime class
 - exitCall method, 2-46
 - getCallerClass method, 2-35
- Orb Class
 - init method, 2-11
- ORBClass property, 2-13
- ORBSingletonClass property, 2-13
- output
 - redirecting, 3-24

P

- packages
 - DBMS_JAVA, 4-3
 - oracle.jdbc2, 2-8
 - protected, 5-25
- performance, 1-15, 6-1 to 6-17
- Permission class, 5-7, 5-13, 5-14, 5-19
- permissions, 4-5, 5-2 to 5-25
 - administrating, 5-12
 - assigning, 5-4, 5-6
 - creating, 5-14
 - deleting, 5-18
 - disabling, 5-17
 - enabling, 5-17
 - FilePermission, 6-6
 - granting, 5-5, 5-8, 5-10
 - granting policy, 5-12
 - grouped into roles, 5-24
 - JAVA_ADMIN role, 5-21
 - JAVA_DEPLOY role, 6-6
 - JAVADEBUGPRIV role, 5-24
 - JAVASYSPRIV role, 5-23
 - JAVAUSERPRIV role, 5-23
 - PUBLIC, 5-22
 - restricting, 5-5, 5-9, 5-10
 - specifying policy, 5-4
 - SYS permissions, 5-22
 - types, 5-7, 5-19
- policy table
 - managing, 5-12
 - modifying, 5-6
 - setting permissions, 5-6
 - viewing, 5-6
- PolicyTable class
 - specifying policy, 5-4
 - updating, 5-5, 5-14
- PolicyTableManager class
 - delete method, 5-18
 - disable method, 5-18
 - enable method, 5-18
 - revoke method, 5-18
- PolicyTablePermission, 5-8, 5-12, 5-19, 5-20, 5-21, 5-22
- polymorphism, 1-6

- presentation
 - compatibility for RMI, 3-10
- presentation layer
 - GIOP, 4-8
- privileges
 - database, 5-2
- .properties files, 2-14, 2-23, 2-25
- property
 - ORBClass, 2-13
 - ORBSingletonClass, 2-13
- PropertyPermission, 5-8, 5-19, 5-21, 5-22, 5-24
- PUBLIC permissions, 5-22
- publish tool, 2-31
- publishing, 2-7, 2-15, 2-30, 3-2
 - example, 3-5

R

- ReflectPermission, 5-8, 5-19, 5-21, 5-22
- registerCallback method, 6-13
- requirements
 - JDK version, 4-8
- reset_compiler_option method, 2-17, 4-4
- resolver, 2-19 to 2-23
 - default, 2-20
 - defined, 2-14, 2-15, 2-20, 2-33, 3-2
 - example, 3-4
 - ignoring non-existent references, 2-20, 2-22
- resource schema object, 2-14, 2-23, 2-25
- restart method, 3-22
- restart_debugging method, 3-22, 4-5
- restrict method, 5-9
- restrict_permission method, 4-5, 5-9, 5-10
- revoke method, 5-18
- revoke_permission method, 4-5, 5-18
- RMI
 - support, 3-10
- RuntimePermission, 5-8, 5-19, 5-21, 5-22, 5-23

S

- schema object
 - defined, 2-23
 - name, 2-32
 - using, 2-14

- security, 5-1 to 5-26
 - book recommendations, 5-4
 - CORBA, 5-2
 - EJB, 5-2
 - Java 2, 5-3
 - JDBC, 5-2
 - JVM, 4-5
 - network, 5-2
 - operating system resources, 2-39
- SecurityManager class, 5-5
- SecurityPermission, 5-8, 5-19, 5-21, 5-22
- .ser files, 2-14, 2-23, 2-25
- SerializablePermission, 5-8, 5-19, 5-21, 5-23
- serialization, 2-36
- ServerSocket class, 2-43
- Servlets, 3-2
- session
 - coordination with JVM, 2-4
 - definition, 2-2
 - footprint, 1-14
 - lifetime, 2-5
 - role in Java execution, 2-3
 - timeout, 2-5
- session shell tool, 3-9
- set_compiler_option method, 2-17, 4-4
- set_compiler_option procedure, 3-19
- set_output method, 3-24, 4-4
- SHARED_POOL_SIZE parameter
 - default, 4-6
 - defined, 6-8
 - errors, 6-11
 - minimum value, 4-3
- shortname method, 2-29, 2-33, 4-3
- Socket class, 2-43
- SocketPermission, 5-8, 5-19, 5-21, 5-23, 5-24
- sockets
 - across calls, 2-38, 2-43
 - defined, 2-43
 - lifetime, 2-39, 2-43
- source schema object, 2-14, 2-23, 2-24
- SQL
 - query, 3-2, 3-11
- SQLJ
 - accessing SQL, 1-18
 - converting, 3-17

- defined, xii, 1-17, 1-19, 3-2, 3-11
- documentation, xii, 1-22
- example, 3-13
- interoperates with PL/SQL, 3-17
- running, 3-16
- translating, 3-16
- typing paradigm, 3-15
- using JDBC, 1-19
- .sqlj files, 2-14, 2-23, 2-24
- sqlj utility, 3-16
- SSL, 4-8
 - configuring, 4-6
- start method, 3-22
- start_debugging method, 3-21, 4-5
- static variable, 2-5
 - end of call migration, 6-12
- stop method, 3-22
- stop_debugging method, 3-21, 4-5
- SYS
 - assigned permissions, 5-22
 - security permissions, 5-19
- System class
 - getProperty method, 3-24

T

- ThreadDeathException, 2-46
- threading
 - applications, 2-44
 - lifecycle, 2-45
 - model, 1-13, 2-44
 - using in JServer, 2-38
- timeout, 2-5
- trigger
 - using Java stored procedures, 3-3

U

- user interface, 2-31
- USER_ERRORS view, 2-16
- USER_JAVA_POLICY view, 5-6, 5-20
- USER_OBJECTS view, 2-25, 2-28, 4-4

V

- VSSGASTAT table, 6-10
- variables
 - static, 2-5
- version
 - retrieving, 3-24
 - Visigenic, 2-10
- Visigenic
 - version supported, 2-10

W

- web sites, xii