

# Oracle® TimesTen In-Memory Database

## Java Developer's Guide



Release 26.1  
F54454-02  
March 2026

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 1996, 2026, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

# Contents

## 1 Java Development Environment

---

Installing TimesTen and Supported Java Components	1
Setting the Environment for Java Development	1
Setting the Module Path or Class Path for Java Development	2
About TimesTen Quick Start and Sample Applications	5

## 2 Working with TimesTen Databases in JDBC

---

Prerequisites for Java Development with TimesTen	1
Key JDBC Classes and Interfaces	1
Package Imports	2
Support for Interfaces in the java.sql Package	3
Support for Classes in the java.sql Package	5
Support Notes for Classes in the java.io Package	5
Support for Interfaces and Classes in the javax.sql Package	6
TimesTen JDBC Extensions	7
Additional TimesTen Classes and Interfaces	8
Management of TimesTen Database Connections	9
About TimesTen DSNs	9
Using Java Wrapper Functionality for Connections	9
Creating a Connection URL for the Database and Specifying Connection Attributes	10
Connecting to the Database	10
Disconnecting from the Database	11
Opening and Closing a Direct Connection	11
Checking Database Validity	12
Database Operations in JDBC	12
Executing Basic SQL Statements	12
Committing or Rolling Back Changes to the Database	14
Setting Autocommit	14
Manually Committing or Rolling Back Changes	14
Using COMMIT and ROLLBACK SQL Statements	15
Managing Multiple Threads	15
Java Escape Syntax and SQL Functions	15
TimesTen Features and Operations in Your Application	16

Using Java Wrapper Functionality for Statements	16
Working with TimesTen Result Sets: Hints and Restrictions	16
Fetching Multiple Rows of Data	17
Optimizing Query Performance	18
Parameter Binding and Statement Execution	19
Preparing SQL Statements and Setting Input Parameters	19
Working with Output and Input/Output Parameters	23
Binding Duplicate Parameters in SQL Statements	24
Binding Duplicate Parameters in PL/SQL	25
Working with Associative Arrays	26
Overview of Associative Arrays	26
TimesTen Methods for Associative Arrays	27
Type Mapping for Associative Arrays	29
Binding Associative Arrays	29
Working with REF CURSORS	30
Working with DML Returning (RETURNING INTO Clause)	32
Working with Rowids	34
Large Objects (LOBs)	34
About LOBs	35
LOB Objects in JDBC	35
Differences Between TimesTen LOBs and Oracle Database LOBs	36
LOB Factory Methods	36
LOB Getter and Setter Methods	37
TimesTen LOB Interface Methods	37
LOB Prefetching	38
Passthrough LOBs	39
Using CALL to Execute Procedures and Functions	39
Setting a Timeout or Threshold for Executing SQL Statements	41
Setting a Timeout Duration for SQL Statements	41
Setting a Threshold Duration for SQL Statements	42
Configuring the Result Set Buffer Size in Client/Server Using JDBC	42
Features for Use with Cache	43
Setting the Oracle Database Password	44
Setting Temporary Passthrough Level with the ttOptSetFlag Built-In Procedure	44
Determining Passthrough Status	44
Managing Cache Groups	45
Features for Use with Replication	45
Error Handling	45
Error and Warning Levels	46
Fatal Errors	46
Non-Fatal Errors	46
Warnings	47

Abnormal Termination	47
Reporting Errors and Warnings	47
Catching and Responding to Specific Errors	48
Rolling Back Failed Transactions	49
Retrying After Transient Errors (JDBC)	50
JDBC Support for Automatic Client Failover	52
About Automatic Client Failover	52
Features and Functionality of JDBC Support for Automatic Client Failover	53
General Client Failover Features	53
Client Failover Features for Pooled Connections	54
Configuration of Automatic Client Failover	54
Synchronous Detection of Automatic Client Failover in JDBC	55
Asynchronous Detection of Automatic Client Failover in JDBC	56
Implement a Client Failover Event Listener	56
Register the Client Failover Listener Instance	57
Remove the Client Failover Listener Instance	58
JDBC Application Action in the Event of Failover	58
Application Steps for Failover	58
Failover Delay and Retry Settings	59

### 3 Using JMS/XLA for Event Management

---

JMS/XLA Concepts and Features	1
JMS/XLA Concepts	1
How XLA Reads Records from the Transaction Log	2
XLA and Materialized Views	3
XLA Bookmarks	4
How Bookmarks Work	4
Replicated Bookmarks	5
XLA Bookmarks and Transaction Log Holds	5
JMS/XLA Configuration File and Topics	6
XLA Updates	7
XLA Update Acknowledgments	8
XLA Acknowledgment Mechanism	8
XLA Acknowledgment Modes	8
Prefetching Updates	9
Acknowledging Updates	9
XLA System Privilege	9
XLA Limitations	9
JMS/XLA and Oracle GDK Dependency	10
JMS/XLA Usage and Functionality	10
Connecting to XLA	10

Monitoring Tables for Updates	11
Receiving and Processing Updates	12
Terminating a JMS/XLA Application	15
Close the Connection	15
Delete Bookmarks	15
Unsubscribe from a Table	16
JMS/XLA as a Replication Mechanism	16
About Using JMS/XLA as a Replication Mechanism	17
Applying JMS/XLA Messages to a Target Database	17
TargetDataStore Error Recovery	18

## 4 Distributed Transaction Processing: JTA

---

Overview of JTA	1
About TimesTen JTA	1
X/Open DTP Model	2
Two-Phase Commit	2
JTA Functionality in TimesTen	3
TimesTen Database Requirements for XA	3
Global Transaction Recovery in TimesTen	3
XA Error Handling in TimesTen	4
Using the JTA API	4
Registering a TimesTen DSN with WebLogic	4
Importing Required Packages	4
Creating a TimesTen XAConnection Object	5
Creating XAResource and Connection Objects	6

## 5 Java Application Tuning

---

Tuning JDBC Applications	1
Use Prepared Statement Pooling	1
Use Arrays of Parameters for Batch Execution	2
Bulk Fetch Rows of TimesTen Data	3
Use the ResultSet Method getString() Sparingly	4
Avoid Data Type Conversions	4
Close Connections, Statements and Result Sets	4
Optimize Queries	5
Tuning JMS/XLA Applications	5
Considerations in Tuning JMS/XLA Applications	5
Configure xlaPrefetch Parameter	5
Reduce Frequency of Update Acknowledgments	5

## 6 JMS/XLA Reference

---

JMS/XLA MapMessage Contents	1
XLA Update Types	1
XLA Flags	3
DML Event Data Formats	4
Table Data	4
Row Data	5
Context Information	5
DDL Event Data Formats	5
CREATE_TABLE	5
DROP_TABLE	6
CREATE_INDEX	6
DROP_INDEX	7
ADD_COLUMNS	7
DROP_COLUMNS	8
CREATE_VIEW	9
DROP_VIEW	9
CREATE_SEQ	10
DROP_SEQ	10
CREATE_SYNONYM	10
DROP_SYNONYM	11
TRUNCATE	11
Data Type Support	11
Data Type Mapping	12
Data Types Character Set	14
JMS Interfaces for Event Handling	14
JMS/XLA Replication API	15
TargetDataStore Interface	15
TargetDataStoreImpl Class	15
JMS Message Header Fields	15

# About This Content

This document covers TimesTen support for JDBC.

## Audience

This guide is for anyone developing or supporting applications that use TimesTen through JDBC.

In addition to familiarity with the particular programming interface you use, you should be familiar with TimesTen, SQL (Structured Query Language), and database operations.

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

## Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

## Related Resources

For reference information on standard JDBC, see the following for information about the `java.sql` and `javax.sql` packages:

<https://docs.oracle.com/javase/8/docs/api/java/sql/package-summary.html>  
<https://docs.oracle.com/javase/8/docs/api/javax/sql/package-summary.html>

Javadoc for standard Java EE classes and interfaces is available at this location:

<https://javaee.github.io/javaee-spec/javadocs/>

Oracle Database documentation is also available on the Oracle documentation website. This may be especially useful for Oracle Database features that TimesTen supports but does not attempt to fully document. In particular, the following Oracle Database documents may be of interest.

- *Oracle Database SQL Language Reference*
- *Oracle Database JDBC Developer's Guide*

## Conventions

The following text conventions are used in this document.

Convention	Meaning
<b>boldface</b>	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

# 1

## Java Development Environment

This chapter provides information about getting started with Java and TimesTen.

- [Installing TimesTen and Supported Java Components](#)
- [Setting the Environment for Java Development](#)
- [Setting the Module Path or Class Path for Java Development](#)
- [About TimesTen Quick Start and Sample Applications](#)

### Installing TimesTen and Supported Java Components

Install and configure TimesTen for your environment.

1. See Overview of the Installation Process in TimesTen Classic in the *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide*.
2. Install the Java Development Kit (JDK) as described in your Java installation documentation.

#### Note

The TimesTen JDBC driver implements the JDBC 4.3 API (Java 9 and up) and is certified to work with Java Runtime Environment (JRE) 11, 17, 21, and 25.

3. If you plan to use the Jakarta Messaging API version 1.1 (Jakarta JMS), download the Jakarta JMS JAR file. From <https://repo1.maven.org/maven2/jakarta/jms/jakarta.jms-api/3.1.0/>, download `jakarta.jms-api-3.1.0.jar`. Older versions of the JAR file aren't certified to work.

You will need to add the path to the JAR file to your module path or class path. See [Setting the Module Path or Class Path for Java Development](#).

### Setting the Environment for Java Development

Before you begin developing Java applications for TimesTen, you must set your environment appropriately. This includes setting environment variables.

Environment variables and runtime access to the Instant Client are configured through the appropriate `ttenv` script in the `timesten_home/bin` directory, where `timesten_home` is the TimesTen instance home directory: `ttenv.sh` and `ttenv.csh` for Linux and UNIX platforms (where which you use depends on your shell) and `ttenv.bat` for Windows platforms.

See Environment Variables and Java Environment Variables in *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide* for information about environment variables, including how to set them using `ttenv` and discussion of the `PATH` and `CLASSPATH` environment variables.

**Note**

TimesTen includes Oracle Instant Client, which is required for certain JDBC features and operations.

## Setting the Module Path or Class Path for Java Development

Compiling any Java application requires the appropriate JARs to be in your module path or class path.

### Java Development Kit

Before running a Java program that loads any of the TimesTen JDBC drivers, configure the JDK library file and path in your Java runtime environment. In Java, class path and module path both tell the JVM and compiler where to find code.

For JDK 11, 17, 21, and 25, the TimesTen JDBC JAR is `timesten_home/install/lib/ttjdbcjdk_version.jar`, where `jdk_version` indicates the JDK version, 11, 17, 21, or 25, for example, `ttjdbc25.jar` for JDK 25. The JAR is also packaged as a Java module with the Java module name, `timesten.jdbc`, so you can use it for module compilation.

The tables below summarize the supported ways to compile and run applications, including both modular and non-modular usage. The options differ slightly between JDK 25 and JDK 11-21.

Table 1-1 JDK 25: Using the TimesTen JDBC Module

Option	Description	Compile Command	Run Command
<b>Option 1: Modular application</b>	Add ttjdbc25.jar to the module path.	<pre>javac --module-path &lt;programe&gt;.java module- info.java ...</pre> <p>where:</p> <ul style="list-style-type: none"> <li>\$module_path is <i>timesten_home/install/lib/ttjdbc25.jar:existing_path_to_modules</i> is the name of the application you want to compile.</li> <li>&lt;programe&gt;.java is the name of the application you want to compile.</li> <li>module-info.java describes the module to include in the application.</li> </ul> <p>The following example shows a sample module-info.java file that requires the timesten.jdbc module.</p> <pre>module my.testmodule {     requires java.sql;     requires timesten.jdbc;     ... }</pre>	<pre>java --module-path \$module_path -- enable-native-access=timesten.jdbc" -- module \$name_of_module/\$classname</pre> <p>where:</p> <ul style="list-style-type: none"> <li>\$module_path is <i>timesten_home/install/lib/ttjdbc25.jar:existing_path_to_modules</i></li> <li>\$name_of_module is the module defined in module-info.java.</li> <li>\$classname is the complete classname including namespace (package name) where the main execution is.</li> </ul>
<b>Option 2: Non-modular or mixed application</b>	If your application is not a module or it uses JARs that are not packaged as Java modules, add the JARs that are modules to the module path and continue to use classpath for the rest of the non-module dependencies.	<pre>javac -cp \$CLASSPATH -- module-path \$module_path ...</pre> <p>where \$module_path is <i>timesten_home/install/lib/ttjdbc25.jar:existing_path_to_modules</i>.</p>	<pre>java -cp \$CLASSPATH --module- path \$module_path --enable-native- access="timesten.jdbc" ...</pre>

**Note**

Every time you want to compile a modular application the classes must have a package name (\$classname).

**Table 1-2 JDK 11, 17, 21: Using the TimesTen JDBC Module or Classpath**

Option	Description	Compile Command	Run Command
<b>Option 1: Modular application</b>	Add <code>ttjdbcjdk_version.jar</code> to the module path.	<code>javac --module-path \$module_path ...</code> where <code>\$module_path</code> is <code>timesten_home/install/lib/ttjdbcjdk_version.jar:existing_path_to_modules</code> .	<code>java --module-path \$module_path --enable-native-access="timesten.jdbc" ...</code> For JDK 11, exclude <code>--enable-native-access</code> : <code>java --module-path \$module_path ...</code>
<b>Option 2: Non-modular or mixed application</b>	If your application is not a module or it uses JARs that are not packaged as Java modules, add the JARs that are modules to the module path and continue to use classpath for the rest of the non-module dependencies.	<code>javac -cp \$CLASSPATH --module-path \$module_path ...</code> where <code>\$module_path</code> is <code>timesten_home/install/lib/ttjdbcjdk_version.jar:existing_path_to_modules</code> .	<code>java -cp \$CLASSPATH --module-path \$module_path --enable-native-access="timesten.jdbc" ...</code>
<b>Option 3: Legacy classpath</b>	Add <code>ttjdbcjdk_version.jar</code> path to the classpath.	<code>javac -cp \$CLASSPATH ...</code> where <code>\$CLASSPATH</code> is <code>timesten_home/install/lib/ttjdbcjdk_version.jar:existing_classpath</code> .	<code>java -cp \$CLASSPATH ...</code>

Java Environment Variables in *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide* discusses the `CLASSPATH` setting for compiling Java applications in TimesTen.

### JMS/XLA Applications

Compiling any JMS/XLA application requires the following to be in your classpath:

JMS Type	Classpath Entries
Jakarta	<code>timesten_home/install/lib/timestenjmsxla.jar</code> <code>timesten_home/install/lib/orai18n.jar</code> In <a href="#">Installing TimesTen and Supported Java Components</a> , you downloaded the Jakarta JMS JAR. You need to add the location of the JAR to your classpath: <code>download-directory/jakarta.jms-api-3.1.0.jar</code>
JavaX	<code>timesten_home/install/lib/timestenjmsxla.jar</code> <code>timesten_home/install/lib/orai18n.jar</code> <code>timesten_home/install/3rdparty/jms1.1/lib/jms.jar</code>

`timestenjmsxla.jar` is for use with both Jakarta JMS and JavaX JMS, but note `download-directory/jakarta.jms-api-3.1.0.jar` in the classpath determines Jakarta JMS and `download-directory/install/3rdparty/jms1.1/lib/jms.jar` in the classpath determines JavaX JMS. If both JARs are in the classpath, TimesTen defaults to Jakarta JMS.

When running your modular JMS/XLA application, you can add the above `timestenjmsxla.jar` path to the module path and add the module name `timesten.jdbc` to `--enable-native-access` option.

```
java --module-path $module_path --enable-native-access="timesten.jdbc,timesten.jmsxla"
```

where `$module_path` is the module path that includes `timesten_home/install/lib/ttjdbcjdk_version.jar` and `timesten_home/install/lib/timestenjmsxla.jar:existing_path_to_modules`.

### Classpath Entry for `ttenv`

Use the appropriate `ttenv` script to set your environment, as discussed in [Setting the Environment for Java Development](#).

#### Note

- For each TimesTen instance, the `timesten_home/install` path is a symbolic link to `installation_dir`, where TimesTen is installed.
- On Windows, there is only one TimesTen instance per installation, and `timesten_home` refers to `installation_dir\instance`.

## About TimesTen Quick Start and Sample Applications

TimesTen Quick Start and sample applications are available from the TimesTen GitHub location. It includes tutorials, how-to instructions, and sample applications. After you have configured your environment, you can confirm that everything is set up correctly by compiling and running the sample applications. TimesTen applications are located under the Quick Start `sample_code` directory. For instructions on compiling and running them, see the instructions in the subdirectories.

The following content is included:

- Schema and setup: The `build_sampleldb` script (`.sh` on Linux or UNIX or `.bat` on Windows) creates a sample database and schema. Run this script before using the sample applications.
- Environment and setup: The `ttquickstartenv` script (`.sh` or `.csh` on Linux or UNIX, `.bat` on Windows, or as applicable for your system), a superset of the `ttenv` script typically used for TimesTen setup, sets up the environment. Run this script each time you enter a session where you want to compile or run any of the sample applications.
- Sample applications and setup: The Quick Start provides sample applications and their source code for JDBC.

# 2

## Working with TimesTen Databases in JDBC

This chapter describes the basic procedures for writing a Java application to access a TimesTen database.

The following topics are covered:

- [Prerequisites for Java Development with TimesTen](#)
- [Key JDBC Classes and Interfaces](#)
- [Management of TimesTen Database Connections](#)
- [Database Operations in JDBC](#)
- [TimesTen Features and Operations in Your Application](#)
- [Error Handling](#)
- [JDBC Support for Automatic Client Failover](#)

### Prerequisites for Java Development with TimesTen

Before attempting to write a TimesTen Java application, you need to complete certain prerequisite tasks.

Prerequisite task	What you do
Create a database.	Follow the procedures described in Managing TimesTen Databases in <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
Configure the Java environment.	Follow the procedures described in <a href="#">Setting the Environment for Java Development</a> .
Compile and execute the TimesTen Java sample applications.	Refer to <a href="#">About TimesTen Quick Start and Sample Applications</a> .

After you have successfully executed the TimesTen Java sample applications, your development environment is set up correctly and ready for you to create applications that access a database.

### Key JDBC Classes and Interfaces

There is support for standard and TimesTen-specific JDBC packages, classes, and interfaces. TimesTen supports JDBC 4.3 API (Java 9 and up).

The following topics are covered:

- [Package Imports](#)
- [Support for Interfaces in the java.sql Package](#)
- [Support for Classes in the java.sql Package](#)
- [Support Notes for Classes in the java.io Package](#)

- [Support for Interfaces and Classes in the javax.sql Package](#)
- [TimesTen JDBC Extensions](#)
- [Additional TimesTen Classes and Interfaces](#)

For reference information on standard JDBC, see the following for information about the `java.sql` and `javax.sql` packages:

JDK Version	Documentatoin for java.sql and java.sql Packages
25	<a href="https://docs.oracle.com/en/java/javase/25/docs/api/java.sql/java/sql/package-summary.html">https://docs.oracle.com/en/java/javase/25/docs/api/java.sql/java/sql/package-summary.html</a> <a href="https://docs.oracle.com/en/java/javase/25/docs/api/java.sql/javax/sql/package-summary.html">https://docs.oracle.com/en/java/javase/25/docs/api/java.sql/javax/sql/package-summary.html</a>
21	<a href="https://docs.oracle.com/en/java/javase/21/docs/api/java.sql/java/sql/package-summary.html">https://docs.oracle.com/en/java/javase/21/docs/api/java.sql/java/sql/package-summary.html</a> <a href="https://docs.oracle.com/en/java/javase/21/docs/api/java.sql/javax/sql/package-summary.html">https://docs.oracle.com/en/java/javase/21/docs/api/java.sql/javax/sql/package-summary.html</a>
17	<a href="https://docs.oracle.com/en/java/javase/17/docs/api/java.sql/java/sql/package-summary.html">https://docs.oracle.com/en/java/javase/17/docs/api/java.sql/java/sql/package-summary.html</a> <a href="https://docs.oracle.com/en/java/javase/17/docs/api/java.sql/javax/sql/package-summary.html">https://docs.oracle.com/en/java/javase/17/docs/api/java.sql/javax/sql/package-summary.html</a>
11	<a href="https://docs.oracle.com/en/java/javase/11/docs/api/java.sql/java/sql/package-summary.html">https://docs.oracle.com/en/java/javase/11/docs/api/java.sql/java/sql/package-summary.html</a> <a href="https://docs.oracle.com/en/java/javase/11/docs/api/java.sql/javax/sql/package-summary.html">https://docs.oracle.com/en/java/javase/11/docs/api/java.sql/javax/sql/package-summary.html</a>

For reference information on TimesTen JDBC extensions, refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## Package Imports

You will need to import some JDBC packages for your application.

Import the standard JDBC package in any Java program that uses JDBC:

```
import java.sql.*;
```

If you are going to use data sources or pooled connections, also import the standard extended JDBC package:

```
import javax.sql.*;
```

Import the TimesTen JDBC package:

```
import com.timesten.jdbc.*;
```

To use XA data sources for JTA, also import this TimesTen package:

```
import com.timesten.jdbc.xa.*;
```

## Support for Interfaces in the java.sql Package

TimesTen supports the `java.sql` interfaces.

The supported interfaces are indicated in [Table 2-1](#), with TimesTen-specific support and restrictions noted.

Also see [TimesTen JDBC Extensions](#).

**Table 2-1 Supported java.sql Interfaces**

Interface in java.sql	Remarks on TimesTen Support
Blob	<ul style="list-style-type: none"> <li>The <code>position()</code> method, which returns the position where a specified byte pattern or BLOB pattern begins, is not supported.</li> </ul>
CallableStatement	<ul style="list-style-type: none"> <li>You must pass parameters to <code>CallableStatement</code> by position, not by name.</li> <li>You cannot use SQL escape syntax.</li> <li>There is no support for <code>Array</code>, <code>Struct</code>, or <code>Ref</code>.</li> <li>There is no support for <code>Calendar</code> for <code>setDate()</code>, <code>getDate()</code>, <code>setTime()</code>, or <code>getTime()</code>.</li> </ul>
Clob	<ul style="list-style-type: none"> <li>The <code>position()</code> method, which returns the position where a specified character pattern or CLOB pattern begins, is not supported.</li> </ul>
Connection	<ul style="list-style-type: none"> <li>There is no support for savepoints.</li> <li>TimesTen supports Read Committed and Serializable isolation levels:</li> </ul> <pre>setTransactionIsolation     (TRANSACTION_READ_COMMITTED);  setTransactionIsolation     (TRANSACTION_SERIALIZABLE);</pre> <p>See <a href="#">Fetching Multiple Rows of Data</a> for information about the relationship between prefetching and isolation level. Also see Concurrency Control Through Isolation and Locking in <i>Oracle TimesTen In-Memory Database Operations Guide</i> and Isolation in <i>Oracle TimesTen In-Memory Database Reference</i>.</p> <ul style="list-style-type: none"> <li>Use <code>beginRequest()</code> and <code>endRequest()</code> as pairs, where <code>beginRequest()</code> disables the auto commit state and <code>endRequest()</code> restores the auto commit state to its previous one.</li> <li>TimesTen does not support sharding keys and methods <code>setShardingKey()</code> and <code>setShardingKeyIfValid()</code>.</li> </ul>
ConnectionBuilder	<ul style="list-style-type: none"> <li>Created by <code>DataSource</code> object using <code>createConnectionBuilder()</code>.</li> <li>TimesTen does not support <code>ShardingKey()</code> and <code>superShardingKey()</code>.</li> </ul>
DatabaseMetaData	<ul style="list-style-type: none"> <li>There are no restrictions.</li> <li>The <code>supportsRefCursors()</code> method returns <code>TRUE</code>.</li> <li>The <code>supportsSharding()</code> method returns <code>FALSE</code>.</li> <li>The <code>getMaxLogicalLobSize()</code> method returns the maximum number of bytes that TimesTen allows for the logical size of a LOB.</li> </ul>

**Table 2-1 (Cont.) Supported java.sql Interfaces**

Interface in java.sql	Remarks on TimesTen Support
Driver	TimesTen does not use <code>java.util.logging</code> , so the <code>getParentLogger()</code> method returns <code>SQLFeatureNotSupportedException</code> . (The TimesTen driver preceded <code>java.util.logging</code> functionality and uses its own logging mechanism.)
NClob	The <code>position()</code> method, which returns the position where a specified character pattern or NCLOB pattern begins, is not supported.
ParameterMetaData	<ul style="list-style-type: none"> <li>The JDBC driver cannot determine whether a column is nullable and always returns <code>parameterNullableUnknown</code> from calls to <code>isNullable()</code>.</li> <li>The <code>getScale()</code> method returns 1 for <code>VARCHAR</code>, <code>NVARCHAR</code>, and <code>VARBINARY</code> data types if they are <code>INLINE</code>. (Scale is of no significance to these data types.)</li> </ul>
PreparedStatement	<ul style="list-style-type: none"> <li>There is no support for <code>getMetaData()</code> in <code>PreparedStatement</code>.</li> <li>There is no support for <code>Array</code>, <code>Struct</code>, or <code>Ref</code>.</li> <li>Settings using <code>setObject(java.util.Calendar)</code> and <code>setDate(java.util.Date)</code> are mapped to <code>TIMESTAMP</code>. There is no support for the <code>Calendar</code> type in <code>setDate()</code>, <code>getDate()</code>, <code>setTime()</code>, or <code>getTime()</code>.</li> </ul>
ResultSet	<ul style="list-style-type: none"> <li>There is support for <code>getMetaData()</code> in <code>ResultSet</code>.</li> <li>You cannot have multiple open <code>ResultSet</code> objects per statement.</li> <li>You cannot specify the holdability of a result set, so a cursor cannot remain open after it has been committed.</li> <li>There is no support for scrollable or updatable result sets.</li> <li>There is no support for <code>Array</code>, <code>Struct</code>, or <code>Ref</code>.</li> <li>There is no support for the <code>Calendar</code> type in <code>setDate()</code>, <code>getDate()</code>, <code>setTime()</code>, or <code>getTime()</code>.</li> <li>In the TimesTen implementation, <code>setFetchSize()</code> and <code>getFetchSize()</code> allow you to set or get the maximum number of rows of data in a client/server result set buffer. See <a href="#">Configuring the Result Set Buffer Size in Client/Server Using JDBC</a></li> <li>See <a href="#">Working with TimesTen Result Sets: Hints and Restrictions</a>.</li> </ul>
ResultSetMetaData	<ul style="list-style-type: none"> <li>The <code>getPrecision()</code> method returns 0 for undefined precision.</li> <li>The <code>getScale()</code> method returns -127 for undefined scale.</li> </ul>
RowId	<ul style="list-style-type: none"> <li>The <code>ROWID</code> data type can be accessed using the <code>RowId</code> interface.</li> <li>Output and input/output rowids can be registered as <code>Types.ROWID</code>.</li> <li>Metadata methods return <code>Types.ROWID</code> and <code>RowId</code> as applicable.</li> </ul>
ShardingKey and ShardingKeyBuilder	TimesTen does not support sharding keys interfaces and methods.

**Table 2-1 (Cont.) Supported java.sql Interfaces**

Interface in java.sql	Remarks on TimesTen Support
Statement	<ul style="list-style-type: none"> <li>TimesTen does not support auto-generated keys.</li> <li>In TimesTen, the <code>cancel()</code> method delegates to the ODBC function <code>SQLCancel</code>. For details about the TimesTen implementation of this function, see ODBC 2.5 Function Support in <i>Oracle TimesTen In-Memory Database C Developer's Guide</i>. The function is also supported for ODBC 3.5.</li> <li>See <a href="#">Managing Cache Groups</a> for special TimesTen functionality of the <code>getUpdateCount()</code> method with cache groups.</li> <li>In the TimesTen implementation, <code>setFetchSize()</code> and <code>getFetchSize()</code> allow you to set or get the maximum number of rows of data in a client/server result set buffer. See <a href="#">Configuring the Result Set Buffer Size in Client/Server Using JDBC</a></li> </ul>
Wrapper	<ul style="list-style-type: none"> <li>TimesTen exposes <code>TimesTenConnection</code>, <code>TimesTenCallableStatement</code>, <code>TimesTenPreparedStatement</code>, and <code>TimesTenStatement</code> through <code>Wrapper</code>.</li> </ul>
XAConnectionBuilder	<ul style="list-style-type: none"> <li>TimesTen does not support <code>password()</code>, <code>shardingKey()</code>, <code>superShardingKey()</code>, and <code>user()</code>.</li> <li>Create <code>TimesTenXAConnectionBuilder</code> with an associated parent <code>TimesTenXADataSource</code> using <code>createXAConnectionBuilder()</code>.</li> </ul>

## Support for Classes in the java.sql Package

TimesTen supports certain `java.sql` classes (and some additional subclasses of `SQLException`).

**Table 2-2 Supported java.sql Classes**

Interface in java.sql	Remarks on TimesTen Support
DataTruncation	No remarks
Date	No remarks
DriverManager	No remarks
DriverPropertyInfo	No remarks
Time	Because TimesTen does not support <code>TIMEZONE</code> in the <code>TIME</code> data type, Java client/server applications should run in the same time zone as the server.
Timestamp	Same consideration for <code>TIMESTAMP</code> as for <code>TIME</code> .
Types	No remarks
SQLException	No remarks
SQLFeatureNotSupportedException	No remarks
SQLWarning	No remarks

## Support Notes for Classes in the java.io Package

There are special notes for TimesTen support of classes in the `java.io` package.

- An `InputStream` object returned by TimesTen does not support mark or reset operations (specifically, the `mark()`, `markSupported()`, and `reset()` methods).
- The `read()` method of an `InputStream` or `Reader` object returns 0 (zero) if the length of the buffer used in the method call is 0, regardless of the amount of data in the `InputStream` or `Reader` object. Therefore, usage such as the following is problematic if the CLOB length may be 0, such as if it were populated using the `SQL EMPTY_CLOB()` function:

```
java.io.Reader r = myclob.getCharacterStream();
char[] buf = new char[myclob.length()]; //buf for r.read() call
```

Typically, when you call `read()`, -1 is returned if the end of the stream is reached. But in the preceding case, -1 is never returned. Be aware of this when you use streams returned by the BLOB `getBinaryStream()` method, which returns `InputStream`, the CLOB `getAsciiStream()` method, which returns `InputStream`, or the CLOB `getCharacterStream()` method, which returns `Reader`.

## Support for Interfaces and Classes in the `javax.sql` Package

TimesTen supports certain `javax.sql` interfaces:

**Table 2-3 Supported `javax.sql` Interfaces**

Interface in <code>javax.sql</code>	Remarks on TimesTen Support
<code>CommonDataSource</code>	<ul style="list-style-type: none"> <li>• Implemented by <code>TimesTenDataSource</code>.</li> <li>• TimesTen does not use <code>java.util.logging</code>, so the <code>getParentLogger()</code> method, specified in <code>CommonDataSource</code>, returns <code>SQLException</code>. (The TimesTen driver preceded <code>java.util.logging</code> functionality and uses its own logging mechanism.)</li> </ul>
<code>ConnectionPoolDataSource</code>	Implemented by <code>ObservableConnectionDS</code> .
<code>DataSource</code>	<ul style="list-style-type: none"> <li>• Implemented by <code>TimesTenDataSource</code>.</li> <li>• TimesTen does not support <code>getParentLogger()</code>, <code>createShardingKeyBuilder()</code>, and <code>unwrap()</code>.</li> <li>• The <code>isWrapperFor()</code> method returns <code>FALSE</code>.</li> <li>• See <a href="https://github.com/oracle-samples/oracle-timesten-samples/blob/master/quickstart/classic/sample_code/jdbc/level2.java">https://github.com/oracle-samples/oracle-timesten-samples/blob/master/quickstart/classic/sample_code/jdbc/level2.java</a> for an example of <code>DataSource</code> implementation.</li> </ul>
<code>PooledConnection</code>	implemented by <code>ObservableConnection</code> .
<code>XADataSource</code>	<ul style="list-style-type: none"> <li>• Implemented by <code>TimesTenXADataSource</code> (in package <code>com.timesten.jdbc.xa</code>).</li> <li>• TimesTen does not support <code>getParentLogger()</code> and <code>createShardingKeyBuilder()</code></li> </ul>

**Note**

- The TimesTen JDBC driver itself does not implement a database connection pool. The `ObservableConnection` and `ObservableConnectionDS` classes simply implement standard Java EE interfaces, facilitating the creation and management of database connection pools according to the Java EE standard.
- TimesTen supports the `javax.sql` event listener. When using a `PooledConnection` instance, you can register a `ConnectionEventListener` instance to listen for `ConnectionEvent` occurrences.
- You can register a `StatementEventListener` instance in TimesTen; however, `StatementEvent` instances are not supported.

## TimesTen JDBC Extensions

For most scenarios, you can use standard JDBC functionality as supported by TimesTen, but TimesTen also provides extensions in the `com.timesten.jdbc` package for TimesTen-specific features.

This is shown in [Table 2-4](#).

**Table 2-4 TimesTen JDBC Extensions**

Interface	Extends	Remarks
<code>TimesTenBlob</code>	<code>Blob</code>	You can cast <code>Blob</code> instances to <code>TimesTenBlob</code> . This includes features to indicate whether a LOB is an Oracle Database passthrough LOB, free LOB resources, and get a binary stream with position and length specifications. See <a href="#">Large Objects (LOBs)</a> .
<code>TimesTenCallableStatement</code>	<code>CallableStatement</code>	Exposed through <code>java.sql.Wrapper</code> . Supports PL/SQL REF CURSORS. See <a href="#">Working with REF CURSORS</a> . Supports associative array binds with methods to set input parameters and to register and get output parameters. See <a href="#">Working with Associative Arrays</a> .
<code>TimesTenClob</code>	<code>Clob</code>	You can cast <code>Clob</code> instances to <code>TimesTenClob</code> . This includes features to indicate whether a LOB is an Oracle Database passthrough LOB, free LOB resources, and get a character stream with position and length specifications. See <a href="#">Large Objects (LOBs)</a> .

Table 2-4 (Cont.) TimesTen JDBC Extensions

Interface	Extends	Remarks
TimesTenConnection	Connection	Exposed through <code>java.sql.Wrapper</code> .  Provides capabilities such as prefetching rows to improve performance, optimizing query performance, listening to events for automatic client failover, setting the track number for parallel replication schemes where you specify replication tracks, checking database validity, and setting the maximum number of rows or bytes for the client/server result set buffer.  See <a href="#">Fetching Multiple Rows of Data</a> , <a href="#">Optimizing Query Performance</a> , <a href="#">General Client Failover Features</a> , <a href="#">Features for Use with Replication</a> , <a href="#">Checking Database Validity</a> and <a href="#">Configuring the Result Set Buffer Size in Client/Server Using JDBC</a> .
TimesTenNClob	NClob	You can cast <code>NClob</code> instances to <code>TimesTenNClob</code> . This includes features to indicate whether a LOB is an Oracle Database passthrough LOB.  See <a href="#">Large Objects (LOBs)</a> .
TimesTenPreparedStatement	PreparedStatement	Exposed through <code>java.sql.Wrapper</code> .  Supports DML returning. See <a href="#">Working with DML Returning (RETURNING INTO Clause)</a> .  Supports associative array binds with a method to set input parameters. See <a href="#">Working with Associative Arrays</a> .
TimesTenStatement	Statement	Exposed through <code>java.sql.Wrapper</code> .  Provides capabilities for specifying a query threshold or setting or getting the maximum number of rows or bytes for the client/server result set buffer. See <a href="#">Setting a Threshold Duration for SQL Statements</a> and <a href="#">Configuring the Result Set Buffer Size in Client/Server Using JDBC</a> .

## Additional TimesTen Classes and Interfaces

In addition to implementations discussed previously, TimesTen provides interfaces and classes in the `com.timesten.jdbc` package.

Features supported by these interfaces and classes are discussed later in this chapter.

- Use interface `TimesTenTypes` for TimesTen type extensions (such as for TimesTen REF CURSORS).
- Use interface `ClientFailoverEventListener` (and also the `ClientFailoverEvent` class below) for automatic client failover features. See [JDBC Support for Automatic Client Failover](#).
- Use interface `TimesTenVendorCode` for vendor codes used in SQL exceptions.
- Use interface `TimesTenConnectionBuilder` to connect to an optimal element based on a distribution key, element ID, or replica set ID. See [#unique\\_45](#).
- Use class `ClientFailoverEvent` (and also the `ClientFailoverEventListener` interface above) for automatic client failover features.

## Management of TimesTen Database Connections

You can manage database connections with TimesTen.

- [About TimesTen DSNs](#)
- [Using Java Wrapper Functionality for Connections](#)
- [Creating a Connection URL for the Database and Specifying Connection Attributes](#)
- [Connecting to the Database](#)
- [Disconnecting from the Database](#)
- [Opening and Closing a Direct Connection](#)
- [Checking Database Validity](#)

Operations described in this section are based on the `level11` sample application. Refer to [About TimesTen Quick Start and Sample Applications](#).

### About TimesTen DSNs

A DSN (data source name) is a logical name that identifies a TimesTen database and the set of connection attributes user for connecting to the database. The type of DSN you create depends on whether your application connects directly to the database or connects from a client.

For TimesTen Classic, refer to *Oracle TimesTen In-Memory Database Operations Guide*. To connect directly to the database, create a DSN as described in [Creating a DSN on Linux and UNIX for TimesTen](#). To create a client connection to the database, create a DSN as described in [Creating and Configuring Client DSNs on Windows](#) or [Creating and Configuring Client DSNs on Linux and UNIX](#).

After you have created a DSN, your application can connect to the database. The sections that follow describe how to create a JDBC connection to a database using either the JDBC direct driver or the JDBC client driver.

### Using Java Wrapper Functionality for Connections

TimesTen exposes TimesTen-specific implementations through standard `java.sql Wrapper` functionality.

You can use `Wrapper` to retrieve connection objects that implement the `TimesTenConnection` interface and provide access to TimesTen-specific features. The following example returns a

TimesTenConnection object then calls its TimesTen extension `setReplicationTrack()` method.

```
String databaseUrl = null;
...
Connection conn = DriverManager.getConnection(databaseUrl);
If (conn.isWrapperFor(TimesTenConnection.class) ) {
    TimesTenConnection tconn = conn.unwrap(TimesTenConnection.class);
    tconn.setReplicationTrack(4);
}
...
```

## Creating a Connection URL for the Database and Specifying Connection Attributes

To create a JDBC connection, specify a TimesTen connection URL for the database. The format of a TimesTen connection URL is:

```
jdbc:timesten:{direct|client}:dsn=DSNname;[DSNattributes;]
```

The default is `direct`.

For example, the following creates a direct connection to the sample database:

```
String URL = "jdbc:timesten:direct:dsn=sampledb";
```

You can programmatically set or override the connection attributes in the DSN description by specifying attributes in the connection URL.

Refer to [Connection Attributes for Data Manager DSNs or Server DSNs](#) in *Oracle TimesTen In-Memory Database Operations Guide*.

General connection attributes require no special privilege. First connection attributes are set when the database is first loaded, and persist for all connections. Only the instance administrator can load a database with changes to first connection attribute settings. Refer to [Connection Attributes](#) in *Oracle TimesTen In-Memory Database Reference*.

For example, to set the `LockLevel` general connection attribute to 1, create a URL as follows:

```
String URL = "jdbc:timesten:direct:dsn=sampledb;LockLevel=1";
```

## Connecting to the Database

After you have defined a URL, you can use the `getConnection()` method of either `DriverManager` or `TimesTenDataSource` to connect to the database.

If you use the `DriverManager.getConnection()` method, specify the driver URL to connect to the database.

```
import java.sql.*;
...
Connection conn = DriverManager.getConnection(URL);
```

To use the `TimesTenDataSource` method `getConnection()`, first create a data source. Then use the `TimesTenDataSource` method `setUrl()` to set the URL and `getConnection()` to connect:

```
import com.timesten.jdbc.TimesTenDataSource;
import java.sql.*;
...

TimesTenDataSource ds = new TimesTenDataSource();
ds.setUrl("jdbc:timesten:direct:<dsn>");
Connection conn = ds.getConnection();
```

The TimesTen user name and password can be set in the DSN within the URL in the `setUrl()` call, but there are also `TimesTenDataSource` methods to set them separately, as well as to set the Oracle Database password (as applicable):

```
TimesTenDataSource ds = new TimesTenDataSource();
ds.setUser(myttusername); // User name to log in to
TimesTen
ds.setPassword(ttpassword); // Password to log in to TimesTen
ds.setUrl("jdbc:timesten:direct:<dsn>");
ds.setOraclePassword(oraclepassword); // Password to log in to Oracle
DB
Connection conn = ds.getConnection();
```

Either the `DriverManager.getConnection()` method or the `ds.getConnection()` method returns a `Connection` object (`conn` in this example) that you can use as a handle to the database. See the `level1` sample application for an example on how to use the `DriverManager` method `getConnection()`, and the `level2` and `level3` sample applications for examples of using the `TimesTenDataSource` method `getConnection()`. Refer to [About TimesTen Quick Start and Sample Applications](#).

## Disconnecting from the Database

When you are finished accessing the database, typically call the `Connection` method `close()` to close the connection to the database.

TimesTen connections also support the standard `abort()` method, as well as standard try-with-resource functionality using `java.lang.AutoCloseable`.

If an error has occurred, you may want to roll back the transaction before disconnecting from the database. See [Non-Fatal Errors](#) and [Rolling Back Failed Transactions](#).

## Opening and Closing a Direct Connection

This example shows the general framework for an application that uses the `DriverManager` class to create a direct connection to the sample database, execute some SQL, and then close the connection.

See the `level1` sample application for a working example. (See [About TimesTen Quick Start and Sample Applications](#) regarding the sample applications.)

```
String URL = "jdbc:timesten:dsn=sampled";
Connection conn = null;
```

```
try {
    Class.forName("com.timesten.jdbc.TimesTenDriver");
} catch (ClassNotFoundException ex) {
    // See Error Handling
}

try {
    // Open a connection to TimesTen
    conn = DriverManager.getConnection(URL);

    // Report any SQLWarnings on the connection
    // See Reporting Errors and Warnings

    // Do SQL operations
    // See TimesTen Features and Operations in Your Application below

    // Close the connection to TimesTen
    conn.close();

    // Handle any errors
} catch (SQLException ex) {
    // See Error Handling
}
```

## Checking Database Validity

Applications can call this `TimesTenConnection` method to detect whether the database is valid.

```
boolean isDataStoreValid() throws java.sql.SQLException
```

It returns `true` if the database is valid, or `false` if the database is in an invalid state, such as due to system or application failure.

## Database Operations in JDBC

There are some general aspects of database operations in JDBC.

- [Executing Basic SQL Statements](#)
- [Committing or Rolling Back Changes to the Database](#)
- [Managing Multiple Threads](#)
- [Java Escape Syntax and SQL Functions](#)

## Executing Basic SQL Statements

You can use the `createStatement()` method of a `Connection` instance, and the `executeUpdate()` or `executeQuery()` method of a `Statement` instance, to execute a SQL statement within a Java application.

See *Working with Data in a TimesTen Database* in *Oracle TimesTen In-Memory Database Operations Guide*.

Unless statements are prepared in advance, use the execution methods of a `Statement` object, such as `execute()`, `executeUpdate()` or `executeQuery()`, depending on the nature of the SQL statement and any returned result set.

For SQL statements that are prepared in advance, use the same execution methods of a `PreparedStatement` object.

The `execute()` method returns `true` if there is a result set (for example, on a `SELECT`) or `false` if there is no result set (for example, on an `INSERT`, `UPDATE`, or `DELETE`). The `executeUpdate()` method returns the number of rows affected. For example, when executing an `INSERT` statement, the `executeUpdate()` method returns the number of rows inserted. The `executeQuery()` method returns a result set, so it should only be called when a result set is expected (for example, when executing a `SELECT` statement).

### Note

- Typically, use the `Statement`, `PreparedStatement`, or `CallableStatement` method `close()` to close a statement you have finished using. TimesTen statements also support standard try-with-resource functionality using `java.lang.AutoCloseable`.
- See [Working with TimesTen Result Sets: Hints and Restrictions](#).

This example uses the `executeUpdate()` method on the `Statement` object to execute an `INSERT` statement to insert data into the `customer` table in the current schema. The connection must have been opened, which is not shown.

```
Connection conn = null;
Statement stmt = null;
...
// [Code to open connection. See Connecting to the Database...]
...
try {
    stmt = conn.createStatement();
    int numRows = stmt.executeUpdate("insert into customer values"
        + "(40, 'West', 'Big Dish', '123 Signal St.');"
    }
catch (SQLException ex) {
    ...
}
```

The following example uses an `executeQuery()` call on the `Statement` object to execute a `SELECT` statement on the `customer` table in the current schema and display the returned `java.sql.ResultSet` instance:

```
Statement stmt = null;
. . . . .
try {
    ResultSet rs = stmt.executeQuery("select cust_num, region, " +
        "name, address from customer");
    System.out.println("Fetching result set...");
    while (rs.next()) {
        System.out.println("\n Customer number: " + rs.getInt(1));
        System.out.println(" Region: " + rs.getString(2));
    }
}
```

```
        System.out.println(" Name: " + rs.getString(3));
        System.out.println(" Address: " + rs.getString(4));
    }
}
catch (SQLException ex) {
    ex.printStackTrace();
}
```

## Committing or Rolling Back Changes to the Database

This section discusses autocommit and manual commits or rollbacks, assuming a JDBC Connection object `myconn` and Statement object `mystmt`. The following topics are covered:

- [Setting Autocommit](#)
- [Manually Committing or Rolling Back Changes](#)
- [Using COMMIT and ROLLBACK SQL Statements](#)

### Note

All open cursors on the connection are closed upon transaction commit or rollback in TimesTen.

Refer to Transaction Overview in *Oracle TimesTen In-Memory Database Operations Guide*.

## Setting Autocommit

A TimesTen connection has autocommit enabled by default, but for performance reasons it is recommended that you disable it. You can use the Connection method `setAutoCommit()` to enable or disable autocommit.

Disable autocommit as follows:

```
myconn.setAutoCommit(false);
// Report any SQLWarnings on the connection.
// See Reporting Errors and Warnings.
```

### Note

A `setAutoCommit()` call results in a commit only when the call actually changes the autocommit setting.

## Manually Committing or Rolling Back Changes

If autocommit is disabled, you must use the Connection method `commit()` to manually commit transactions, or the `rollback()` method to roll back changes. Consider the following example.

```
myconn.commit();
```

Or:

```
myconn.rollback();
```

## Using COMMIT and ROLLBACK SQL Statements

You can prepare and execute `COMMIT` and `ROLLBACK` SQL statements the same way as other SQL statements. Using `COMMIT` and `ROLLBACK` statements has the same effect as using the Connection methods `commit()` and `rollback()`.

For example:

```
mystmt.execute("COMMIT");
```

## Managing Multiple Threads

When your application has a direct connection to the database, TimesTen functions share stack space with your application. In multithreaded environments it is important to avoid overrunning the stack allocated to each thread, as this can cause a program to fail in unpredictable ways that are difficult to debug. The amount of stack space consumed by TimesTen calls varies depending on the SQL functionality used. Most applications should set thread stack space between 34 KB and 72 KB.

The amount of stack space allocated for each thread is specified by the operating system when threads are created. On Windows, you can use the TimesTen debug driver and link your application against the Visual C++ debug C library to enable stack probes that raise an identifiable exception if a thread attempts to grow its stack beyond the amount allocated.

### Note

In multithreaded applications, a thread that issues requests on different connection handles to the same database may encounter lock conflicts with itself. TimesTen returns lock timeout and deadlock errors in this situation.

The `level14` sample application demonstrates the use of multiple threads. Refer to [About TimesTen Quick Start and Sample Applications](#).

## Java Escape Syntax and SQL Functions

When using SQL in JDBC, pay special attention to Java escape syntax. SQL functions such as `UNISTR` use the backslash (`\`) character. You should escape the backslash character.

For example, using the following SQL syntax in a Java application may not produce the intended results:

```
INSERT INTO table1 SELECT UNISTR('\00E4') FROM dual;
```

Escape the backslash character as follows:

```
INSERT INTO table1 SELECT UNISTR('\\00E4') FROM dual;
```

# TimesTen Features and Operations in Your Application

This section provides detailed information on working with data in a TimesTen database.

These topics are discussed here:

- [Using Java Wrapper Functionality for Statements](#)
- [Working with TimesTen Result Sets: Hints and Restrictions](#)
- [Fetching Multiple Rows of Data](#)
- [Optimizing Query Performance](#)
- [Parameter Binding and Statement Execution](#)
- [Working with Associative Arrays](#)
- [Working with REF CURSORS](#)
- [Working with DML Returning \(RETURNING INTO Clause\)](#)
- [Working with Rowids](#)
- [Large Objects \(LOBs\)](#)
- [Using CALL to Execute Procedures and Functions](#)
- [Setting a Timeout or Threshold for Executing SQL Statements](#)
- [Configuring the Result Set Buffer Size in Client/Server Using JDBC](#)
- [Features for Use with Cache](#)
- [Features for Use with Replication](#)

## Using Java Wrapper Functionality for Statements

TimesTen exposes TimesTen-specific implementations through standard `java.sql Wrapper` functionality.

You can use `Wrapper` to retrieve statement objects that implement the `TimesTenStatement`, `TimesTenPreparedStatement`, and `TimesTenCallableStatement` interfaces and provide access to TimesTen-specific features. See [Management of TimesTen Database Connections](#) for similar discussion and an example regarding connection objects.

## Working with TimesTen Result Sets: Hints and Restrictions

This section points out what you should know when using `ResultSet` objects from TimesTen.

Use `ResultSet` objects to process query results. In addition, some methods and built-in procedures return TimesTen data in the form of a `ResultSet` object.

### Note

In TimesTen, any operation that ends your transaction, such as a commit or rollback, closes all cursors associated with the connection.

- TimesTen does not support multiple open `ResultSet` objects per statement. TimesTen cannot return multiple `ResultSet` objects from a single `Statement` object without first closing the current result set.
- TimesTen does not support holdable cursors. You cannot specify the holdability of a result set, essentially whether a cursor can remain open after it has been committed.
- `ResultSet` objects are not scrollable or updatable, so you cannot specify `ResultSet.TYPE_SCROLL_SENSITIVE` or `ResultSet.CONCUR_UPDATABLE`.
- Typically, use the `ResultSet` method `close()` to close a result set as soon as you are done with it. For performance reasons, this is especially important for result sets used for both read and update operations and for result sets used in pooled connections.

TimesTen result sets also support standard try-with-resource functionality using `java.lang.AutoCloseable`.

- Calling the `ResultSet` method `getString()` is more costly in terms of performance if the underlying data type is not a string. Because Java strings are immutable, `getString()` must allocate space for a new string each time it is called. Do not use `getString()` to retrieve primitive numeric types, like `byte` or `int`, unless it is absolutely necessary. For example, it is much faster to call `getInt()` on an integer column. Also see [Use the ResultSet Method getString\(\) Sparingly](#).

In addition, for dates and timestamps, the `ResultSet` native methods `getDate()` and `getTimestamp()` have better performance than `getString()`.

- Application performance is affected by the choice of `getXXX()` calls and by any required data transformations after invocation.
- JDBC ignores the setting for the `ConnectionCharacterSet` attribute. It returns data in UTF-16 encoding.

## Fetching Multiple Rows of Data

Fetching multiple rows of data can increase the performance of a client/server application that connects to a database set with Read Committed isolation level.

This section describes the connection-level prefetch implemented in TimesTen.

You can specify the number of rows to be prefetched by calling the `TimesTenConnection` method `setTtPrefetchCount()`. This enables a TimesTen extension that establishes prefetch at the connection level so that all of the statements on the connection use the same prefetch setting..

### Note

The TimesTen prefetch count extension provides no benefit for an application using a direct connection to the database.

When you set the prefetch count to 0, TimesTen uses a default prefetch count according to the isolation level you have set for the database, and sets the prefetch count to that value. With Read Committed isolation level, the default prefetch value is 5. With Serializable isolation level, the default is 128. The default prefetch value is a good setting for most applications. Generally, a higher value may result in better performance for larger result sets, at the expense of slightly higher resource use.

To disable prefetch, set the prefetch count to 1.

Call the `TimesTenConnection` method `getTtPrefetchCount()` to check the current prefetch value.

See the `Connection` interface entry in [Support for Interfaces in the java.sql Package](#). Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

The following example uses a `setTtPrefetchCount()` call to set the prefetch count to 10, then uses a `getTtPrefetchCount()` call to return the prefetch count in the `count` variable.

```
TimesTenConnection conn =
    (TimesTenConnection) DriverManager.getConnection(url);

// set prefetch count to 10 for this connection
conn.setTtPrefetchCount(10);

// Return the prefetch count to the 'count' variable.
int count = conn.getTtPrefetchCount();
```

## Optimizing Query Performance

A TimesTen extension enables applications to optimize read-only query performance in client/server applications by calling the `TimesTenConnection` method `setTtPrefetchClose()` with a setting of `true`.

All transactions should be committed when executed, including read-only transactions. With a `setTtPrefetchClose(true)` call, the server automatically closes the cursor and commits the transaction after the server has prefetched all rows of the result set for a read-only query. This enhances performance by reducing the number of network round-trips between client and server.

The client should still close the result set and commit the transaction, but those calls are executed in the client and do not require a network round trip between the client and server.

### Note

- Do not use multiple statement handles for the same connection with a `setTtPrefetchClose(true)` call. The server may fetch all rows from the result set, commit the transaction, and close the statement handle before the client is finished, resulting in the closing of all statement handles.
- A `true` setting is ignored for TimesTen direct connections and for `SELECT FOR UPDATE` statements.
- Use `getTtPrefetchClose()` to get the current setting (`true` or `false`).

The following example shows usage of `setTtPrefetchClose(true)`.

```
import com.timesten.sql;
...
con = DriverManager.getConnection ("jdbc:timesten:client:" + DSN);
stmt = con.createStatement();
...
con.setTtPrefetchClose(true);
rs = stmt.executeQuery("select * from t");
```

```
while(rs.next())
{
// do the processing
}
rs.close();
con.commit();
```

## Parameter Binding and Statement Execution

This discusses how to bind input or output parameters for SQL statements.

The following topics are covered.

- [Preparing SQL Statements and Setting Input Parameters](#)
- [Working with Output and Input/Output Parameters](#)
- [Binding Duplicate Parameters in SQL Statements](#)
- [Binding Duplicate Parameters in PL/SQL](#)

### Note

The term "bind parameter" as used in TimesTen developer guides (in keeping with ODBC terminology) is equivalent to the term "bind variable" as used in TimesTen PL/SQL documents (in keeping with Oracle Database PL/SQL terminology).

## Preparing SQL Statements and Setting Input Parameters

SQL statements that are to be executed more than once should be prepared in advance by calling the `Connection` method `prepareStatement()`. For maximum performance, prepare parameterized statements.

Be aware of the following:

- The TimesTen binding mechanism (early binding) differs from that of Oracle Database (late binding). TimesTen requires the data types before preparing queries. As a result, there will be an error if the data type of each bind parameter is not specified or cannot be inferred from the SQL statement. This would apply, for example, to the following statement:

```
SELECT 'x' FROM DUAL WHERE ? = ?;
```

You could address the issue as follows, for example.

```
SELECT 'x' from DUAL WHERE CAST(? as VARCHAR2(10)) = CAST(? as
VARCHAR2(10));
```

- By default (when connection attribute `PrivateCommands=0`), TimesTen shares prepared statements between connections, so subsequent prepares of the same statement on different connections execute very quickly.
- Application performance is influenced by the choice of `setXXX()` calls and by any required data transformations before invocation. For example, for time, dates, and timestamps, the `PreparedStatement` native methods `setTime()`, `setDate()` and `setTimestamp()` have better performance than `setString()`.

- For `TT_TINYINT` columns, use `setShort()` or `setInt()` instead of `setByte()` to use the full range of `TT_TINYINT` (0-255).
- Settings using `setObject(java.util.Calendar)` and `setDate(java.util.Date)` are mapped to `TIMESTAMP`.

The following example shows the basics of an `executeQuery()` call on a `PreparedStatement` object. It executes a prepared `SELECT` statement and displays the returned result set.

```
PreparedStatement pSel = conn.prepareStatement("select cust_num, " +
        "region, name, address " +
        "from customer" +
        "where region = ?");

pSel.setInt(1,1);

try {
    ResultSet rs = pSel.executeQuery();

    while (rs.next()) {
        System.out.println("\n Customer number: " + rs.getInt(1));
        System.out.println(" Region: " + rs.getString(2));
        System.out.println(" Name: " + rs.getString(3));
        System.out.println(" Address: " + rs.getString(4));
    }
}
catch (SQLException ex) {
    ex.printStackTrace();
}
```

This next example shows how a single parameterized statement can be substituted for four separate statements.

Rather than execute a similar `INSERT` statement with different values:

```
Statement.execute("insert into t1 values (1, 2)");
Statement.execute("insert into t1 values (3, 4)");
Statement.execute("insert into t1 values (5, 6)");
Statement.execute("insert into t1 values (7, 8)");
```

It is more efficient to prepare a single parameterized `INSERT` statement and use `PreparedStatement` methods `setXXX()` to set the row values before each execute.

```
PreparedStatement pIns = conn.prepareStatement("insert into t1 values
(?,?)");

pIns.setInt(1, 1);
pIns.setInt(2, 2);
pIns.executeUpdate();

pIns.setInt(1, 3);
pIns.setInt(2, 4);
pIns.executeUpdate();

pIns.setInt(1, 5);
pIns.setInt(2, 6);
pIns.executeUpdate();
```

```
pIns.setInt(1, 7);
pIns.setInt(2, 8);
pIns.executeUpdate();

conn.commit();
pIns.close();
```

TimesTen shares prepared statements automatically after they have been committed. For example, if two or more separate connections to the database each prepare the same statement, then the second, third, ... , *n*th prepared statements return very quickly because TimesTen remembers the first prepared statement.

The following example prepares `INSERT` and `SELECT` statements, executes the `INSERT` twice, executes the `SELECT`, and prints the returned result set. For a working example, see the `level1` sample application. (Refer to [About TimesTen Quick Start and Sample Applications](#) regarding the sample applications.)

```
Connection conn = null;
...
// [Code to open connection. See Connecting to the Database...]
...

// Disable auto-commit
conn.setAutoCommit(false);

    // Report any SQLWarnings on the connection
    // See Reporting Errors and Warnings

// Prepare a parameterized INSERT and a SELECT Statement
PreparedStatement pIns =
    conn.prepareStatement("insert into customer values
(?,?,?,?)");

PreparedStatement pSel = conn.prepareStatement
    ("select cust_num, region, name, " +
    "address from customer");

// Data for first INSERT statement
pIns.setInt(1, 100);
pIns.setString(2, "N");
pIns.setString(3, "Fiberifics");
pIns.setString(4, "123 any street");

// Execute the INSERT statement
pIns.executeUpdate();

// Data for second INSERT statement
pIns.setInt(1, 101);
pIns.setString(2, "N");
pIns.setString(3, "Natural Foods Co.");
pIns.setString(4, "5150 Johnson Rd");

// Execute the INSERT statement
pIns.executeUpdate();
```

```
// Commit the inserts
conn.commit();

// Done with INSERTs, so close the prepared statement
pIns.close();

// Report any SQLWarnings on the connection.
reportSQLWarnings(conn.getWarnings());

// Execute the prepared SELECT statement
ResultSet rs = pSel.executeQuery();

System.out.println("Fetching result set...");
while (rs.next()) {
    System.out.println("\n Customer number: " + rs.getInt(1));
    System.out.println(" Region: " + rs.getString(2));
    System.out.println(" Name: " + rs.getString(3));
    System.out.println(" Address: " + rs.getString(4));
}

// Close the result set.
rs.close();

// Commit the select - yes selects must be committed too
conn.commit();

// Close the select statement - we are done with it
pSel.close();
```

The next example prepares three identical parameterized `INSERT` statements for three separate connections. The first prepared `INSERT` for connection `conn1` is shared (inside the TimesTen internal prepared statement cache) with the `conn2` and `conn3` connections, speeding up the prepare operations for `pIns2` and `pIns3`:

```
Connection conn1 = null;
Connection conn2 = null;
Connection conn3 = null;
.....
PreparedStatement pIns1 = conn1.prepareStatement
    ("insert into t1 values (?,?)");

PreparedStatement pIns2 = conn2.prepareStatement
    ("insert into t1 values (?,?)");

PreparedStatement pIns3 = conn3.prepareStatement
    ("insert into t1 values (?,?)");
```

### Note

All optimizer hints, such as join ordering, indexes and locks, must match for the statement to be shared in the internal TimesTen prepared statement cache. Also, if the prepared statement references a temp table, it is only shared within a single connection.

## Working with Output and Input/Output Parameters

You can prepare a statement and set input parameters using `PreparedStatement` methods.

See [Preparing SQL Statements and Setting Input Parameters](#). TimesTen also supports output and input/output parameters, for which you use `java.sql.CallableStatement` instead of `PreparedStatement`, as follows.

1. Use the method `registerOutParameter()` to register an output or input/output parameter, specifying the parameter position (position in the statement) and data type.

This is the standard method as specified in the `CallableStatement` interface:

```
void registerOutParameter(int parameterIndex, int sqlType, int scale)
```

Be aware, however, that if you use this standard version for `CHAR`, `VARCHAR`, `NCHAR`, `NVARCHAR`, `BINARY`, or `VARBINARY` data, TimesTen allocates memory to hold the largest possible value. In many cases this is wasteful.

Instead, you can use the TimesTen extended interface `TimesTenCallableStatement`, which has a `registerOutParameter()` signature that enables you to specify the maximum data length. For `CHAR`, `VARCHAR`, `NCHAR`, and `NVARCHAR`, the unit of length is number of characters. For `BINARY` and `VARBINARY`, it is bytes.

```
void registerOutParameter(int paramIndex,  
                          int sqlType,  
                          int ignore, //This parameter is ignored by  
TimesTen.  
                          int maxLength)
```

2. Use the appropriate `CallableStatement` method `setXXX()`, where `XXX` indicates the data type, to set the input value of an input/output parameter. Specify the parameter position and data value.
3. Use the appropriate `CallableStatement` method `getXXX()` to get the output value of an output or input/output parameter, specifying the parameter position.

### Note

- Check for SQL warnings before processing output parameters. In the event of a warning, output parameters are undefined. See [Error Handling](#).
- You cannot pass parameters to a `CallableStatement` object by name. You must set parameters by position. You cannot use the SQL escape syntax.
- The `registerOutParameter()` signatures specifying the parameter by name are not supported. You must specify the parameter by position.
- SQL structured types are not supported.

This example shows how to use a callable statement with an output parameter. In the `TimesTenCallableStatement` instance, a PL/SQL block calls a function `RAISE_SALARY` that calculates a new salary and returns it as an integer. Assume a `Connection` instance `conn`.

(Refer to Overview of PL/SQL Features in the *Oracle TimesTen In-Memory Database PL/SQL Developer's Guide*.)

```
import java.sql.CallableStatement;
import java.sql.Connection;
import java.sql.Types;
import com.timesten.jdbc.TimesTenCallableStatement;
...
// Prepare to call a PL/SQL stored procedure RAISE_SALARY
CallableStatement cstmt = conn.prepareCall
    ("BEGIN :newSalary := RAISE_SALARY(:name, :inc);
end;");

// Declare that the first param (newSalary) is a return (output) value of
type int
cstmt.registerOutParameter(1, Types.INTEGER);

// Raise Leslie's salary by $2000 (who wanted $3000 but we held firm)
cstmt.setString(2, "LESLIE"); // name argument (type String) is the second
param
cstmt.setInt(3, 2000); // raise argument (type int) is the third param

// Do the raise
cstmt.execute();

// Check warnings. If there are warnings, output parameter values are
undefined.
SQLWarning wn;
boolean warningFlag = false;
if ((wn = cstmt.getWarnings() ) != null) {
    do {
        warningFlag = true;
        System.out.println(wn);
        wn = wn.getNextWarning();
    } while(wn != null);
}

// Get the new salary back
if (!warningFlag) {
    int new_salary = cstmt.getInt(1);
    System.out.println("The new salary is: " + new_salary);
}

// Close the statement and connection
cstmt.close();
conn.close();
...
```

## Binding Duplicate Parameters in SQL Statements

In TimesTen, multiple occurrences of the same parameter name in a SQL statement are considered to be distinct parameters. (This is consistent with Oracle Database support for binding duplicate parameters.)

**Note**

- This discussion applies only to SQL statements issued directly from ODBC, not through PL/SQL, for example.
- "TimesTen mode" for binding duplicate parameters, and the `DuplicateBindMode` connection attribute, are deprecated.

Consider this query:

```
SELECT * FROM employees
WHERE employee_id < :a AND manager_id > :a AND salary < :b;
```

When parameter position numbers are assigned, a number is given to each parameter occurrence without regard to name duplication. The application must, at a minimum, bind a value for the first occurrence of each parameter name. For any subsequent occurrence of a given parameter name, the application has the following choices.

- It can bind a different value for the occurrence.
- It can leave the parameter occurrence unbound, in which case it takes the same value as the first occurrence.

In either case, each occurrence still has a distinct parameter position number.

To use a different value for the second occurrence of `a` in the SQL statement above:

```
pstmt.setXXX(1, ...); /* first occurrence of :a */
pstmt.setXXX(2, ...); /* second occurrence of :a */
pstmt.setXXX(3, ...); /* occurrence of :b */
```

To use the same value for both occurrences of `a`:

```
pstmt.setXXX(1, ...); /* both occurrences of :a */
pstmt.setXXX(3, ...); /* occurrence of :b */
```

Parameter `b` is considered to be in position 3 regardless.

## Binding Duplicate Parameters in PL/SQL

The preceding discussion does not apply to PL/SQL, which has its own semantics. In PL/SQL, you bind a value for each unique parameter name.

An application executing the following block, for example, would bind only one parameter, corresponding to `:a`.

```
DECLARE
  x NUMBER;
  y NUMBER;
BEGIN
  x := :a;
  y := :a;
END;
```

An application executing the following block would also bind only one parameter:

```
BEGIN
  INSERT INTO tab1 VALUES(:a, :a);
END
```

And the same for the following CALL statement:

```
...CALL proc(:a, :a)...
```

An application executing the following block would bind two parameters, with :a as parameter #1 and :b as parameter #2. The second parameter in each INSERT statement would take the same value as the first parameter in the first INSERT statement, as follows.

```
BEGIN
  INSERT INTO tab1 VALUES(:a, :a);
  INSERT INTO tab1 VALUES(:b, :a);
END
```

## Working with Associative Arrays

This section shows how to bind and use associative arrays in TimesTen.

- [Overview of Associative Arrays](#)
- [TimesTen Methods for Associative Arrays](#)
- [Type Mapping for Associative Arrays](#)
- [Binding Associative Arrays](#)

## Overview of Associative Arrays

TimesTen JDBC supports associative arrays, formerly known as index-by tables or PL/SQL tables, as IN, OUT, or IN OUT bind parameters to TimesTen PL/SQL. Associative arrays enable arrays of data to be passed efficiently between a JDBC application and the database.

An associative array is a set of key-value pairs. In TimesTen, for associative array binding (but not for use of associative arrays only within PL/SQL), the keys, or indexes, must be integers (BINARY\_INTEGER or PLS\_INTEGER). The values must be scalar values of the same data type. For example, there could be an array of department managers indexed by department numbers. Indexes are stored in sort order, not creation order.

You can declare an associative array type and then an associative array from PL/SQL as in the following example (note the INDEX BY):

```
declare
  TYPE VARCHARARRTYP IS TABLE OF VARCHAR2(30) INDEX BY BINARY_INTEGER;
  x VARCHARARRTYP;
  ...
```

Also see *Using Associative Arrays from Applications in Oracle TimesTen In-Memory Database PL/SQL Developer's Guide*.

**Note**

- The following types are not supported in binding associative arrays: LOBs, REF CURSORS, TIMESTAMP, ROWID.
- Associative array binding is not allowed in passthrough statements.
- General bulk binding of arrays is not supported in TimesTen JDBC. Varrays and nested tables are not supported as bind parameters.
- Associative array parameters are not supported with JDBC batch execution. (See [Use Arrays of Parameters for Batch Execution](#).)

## TimesTen Methods for Associative Arrays

TimesTen provides extensions through the interfaces `TimesTenPreparedStatement` and `TimesTenCallableStatement` to support associative array binds.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference* for additional information about methods described here.

For an associative array that is a PL/SQL `IN` or `IN OUT` parameter, TimesTen provides the `setPlsqlIndexTable()` method in the `TimesTenPreparedStatement` interface (for an `IN` parameter) and in the `TimesTenCallableStatement` interface (for an `IN OUT` parameter) to set the input associative array.

- `void setPlsqlIndexTable(int paramIndex, java.lang.Object arrayData, int maxLen, int curLen, int elemSqlType, int elemMaxLen)`

Specify the following:

- `paramIndex`: Parameter position within the PL/SQL statement (starting with 1)
- `arrayData`: Array of values to be bound (which can be an array of primitive types such as `int[]` or an array of object types such as `BigDecimal[]`)
- `maxLen`: Maximum number of elements in the associative array (in TimesTen must be same as `curLen`)
- `curLen`: Actual current number of elements in the associative array (in TimesTen must be same as `maxLen`)
- `elemSqlType`: Type of the associative array elements according to `java.sql.Types` (such as `Types.DOUBLE`)
- `elemMaxLen`: For `CHAR`, `VARCHAR`, `BINARY`, or `VARBINARY` associative arrays, the maximum length of each element (in characters for `CHAR` or `VARCHAR` associative arrays, or in bytes for `BINARY` or `VARBINARY` associative arrays)

For example (assuming a `TimesTenPreparedStatement` instance `pstmt`):

```
int maxLen = 3;
int curLen = 3;
// Numeric field can be set with int, float, double types.
// elemMaxLen is set to 0 for numeric types and is ignored.
// elemMaxLen is specified for VARCHAR types.
pstmt.setPlsqlIndexTable
    (1, new int[]{4, 5, 6}, maxLen, curLen, Types.NUMERIC, 0);
pstmt.setPlsqlIndexTable
```

```
(2, new String[]{"Batch1234567890", "2", "3"}, maxLen, curLen,
    Types.VARCHAR, 15);
pstmt.execute();
```

### Note

- The *elemMaxLen* parameter is ignored for types other than CHAR, VARCHAR, BINARY, or VARBINARY. For any of those types, you can use a value of 0 to instruct the driver to set the maximum length of each element based on the actual length of data that is bound. If *elemMaxLen* is set to a positive value, then wherever the actual data length is greater than *elemMaxLen*, the data is truncated to a length of *elemMaxLen*.
- If *curLen* is smaller than the actual number of elements in the associative array, only *curLen* elements are bound.

For an associative array that is a PL/SQL OUT or IN OUT parameter, TimesTen provides two methods in the `TimesTenCallableStatement` interface: `registerIndexTableOutParameter()` to register an output associative array, and `getPlsqlIndexTable()` to retrieve an output associative array. There are two signatures for `getPlsqlIndexTable()`, one to use the JDBC default Java object type given the associative array element SQL type, and one to specify the type.

- `void registerIndexTableOutParameter(int paramIndex, int maxLen, int elemSqlType, int elemMaxLen)`

Specify the following:

- *paramIndex*: Parameter position within the PL/SQL statement (starting with 1)
- *maxLen*: Maximum possible number of elements in the associative array
- *elemSqlType*: Type of the associative array elements according to `java.sql.Types` (such as `Types.DOUBLE`)
- *elemMaxLen*: For CHAR, VARCHAR, BINARY, or VARBINARY associative arrays, the maximum length of each element (in characters for CHAR or VARCHAR associative arrays, or in bytes for BINARY or VARBINARY associative arrays)

### Note

If *elemMaxLen* has a value of 0 or less, the maximum length for the data type is used.

- `java.lang.Object getPlsqlIndexTable(int paramIndex)`

With this method signature, the type of the returned associative array is the JDBC default mapping for the SQL type of the data retrieved. Specify the parameter position within the PL/SQL statement (starting with 1). See [Table 2-5](#).

- `java.lang.Object getPlsqlIndexTable(int paramIndex, java.lang.Class primitiveType)`

With this method signature, in addition to specifying the parameter position, specify the desired type of the returned associative array according to `java.sql.Types` (such as `Types.DOUBLE`). It must be a primitive type.

## Type Mapping for Associative Arrays

When you bind an associative array in Java, match the Java type as closely as possible with the array type for optimal performance. TimesTen does, however, support certain input conversions.

- Strings can be converted to integers or floating point numbers.
- Strings can be converted to `DATE` data if the strings are in TimesTen `DATE` format (`YYYY-MM-DD HH:MI:SS`).

[Table 2-5](#) documents JDBC default mappings for associative array elements.

**Table 2-5 JDBC Default Mappings for Associative Array Elements**

Return type	SQL type
<code>Integer[]</code>	<code>TINYINT, SMALLINT, TT_INTEGER</code>
<code>Long[]</code>	<code>BIGINT</code>
<code>BigDecimal[]</code>	<code>NUMBER</code>
<code>Float[]</code>	<code>BINARY_FLOAT</code>
<code>Double[]</code>	<code>BINARY_DOUBLE</code>
<code>String[]</code>	<code>CHAR, VARCHAR, NCHAR, NVARCHAR</code>
<code>Timestamp[]</code>	<code>DATE</code>

## Binding Associative Arrays

The following code fragment illustrates how to set, register, and retrieve the contents of an `IN OUT` parameter (assuming a connection `conn` and `TimesTenCallableStatement` instance `cstmt`):

```
int maxLen = 3;
int curLen = 3;
anonBlock = "begin AssocArrayEx_inoutproc(:o1); end;";
cstmt = (TimesTenCallableStatement) conn.prepareCall(anonBlock);
cstmt.setPsqlIndexTable
    (1, new Integer[] {1,2,3}, maxLen, curLen, Types.NUMERIC, 0);
cstmt.registerIndexTableOutParameter(1, maxLen, Types.NUMERIC, 0);
cstmt.execute();

int[] ret = (int [])cstmt.getPsqlIndexTable(1, Integer.TYPE);
cstmt.execute();
```

The next example is more complete, showing the mechanism for binding an associative array.

```
TimesTenCallableStatement cstmt = null;
try {
    // Prepare procedure with associative array in parameter
    cstmt = (TimesTenCallableStatement)
```

```

        conn.prepareCall("begin AssociativeArray_proc(:name, :inc);
end;");

// Set up input array and length
String[] name = {"George", "John", "Thomas", "James", "Bill"};
Integer[] salaryInc = {10000, null, 5000, 8000, 9007};
int currentLen = name.length;
int maxLen = currentLen;

// Use elemMaxLen for variable length data types such as
// Types.VARCHAR, Types.CHAR.
int elemMaxLen = 32;

// set input parameter, name as a VARCHAR
cstmt.setPlsqlIndexTable
    (1, name, maxLen, currentLen, Types.VARCHAR, elemMaxLen);
// set input parameter, salaryInc as a number
cstmt.setPlsqlIndexTable
    (2, salaryInc, maxLen, currentLen, Types.NUMERIC, 0);

```

## Working with REF CURSORS

*REF CURSOR* is a PL/SQL concept, a handle to a cursor over a SQL result set that can be passed between PL/SQL and an application. In TimesTen, the cursor can be opened in PL/SQL, then the REF CURSOR can be passed to the application for processing of the result set.

TimesTen supports standard JDBC REF CURSORS as well as TimesTen REF CURSORS.

An application can receive a REF CURSOR *OUT* parameter as follows:

1. Using the `CallableStatement` method `registerOutParameter()`, register the REF CURSOR *OUT* parameter as type `java.sql.Types.REF_CURSOR` (for standard REF CURSORS) or as type `TimesTenTypes.CURSOR` (for TimesTen REF CURSORS). In the `registerOutParameter()` call, specify the parameter position of the REF CURSOR (position in the statement).
2. Retrieve the REF CURSOR using the `CallableStatement` method `getObject()` (for standard REF CURSORS) or the `TimesTenCallableStatement` method `getCursor()` (for TimesTen REF CURSORS), casting the return as `ResultSet`. Specify the parameter position of the REF CURSOR.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference* for additional information about TimesTen JDBC APIs. See PL/SQL REF CURSORS in *Oracle TimesTen In-Memory Database PL/SQL Developer's Guide* for additional information about REF CURSORS.

### Note

For passing REF CURSORS between PL/SQL and an application, TimesTen supports only *OUT* REF CURSORS, from PL/SQL to the application.

The following examples show how to use a callable statement with a TimesTen REF CURSOR then a standard REF CURSOR.

**TimesTen REF CURSOR:**

```
import java.sql.CallableStatement;
import java.sql.Connection;
import java.sql.ResultSet;
import com.timesten.jdbc.TimesTenCallableStatement;
import com.timesten.jdbc.TimesTenTypes;
...
Connection conn = null;
CallableStatement cstmt = null;
ResultSet cursor;
...
// Use a PL/SQL block to open the cursor.
cstmt = conn.prepareCall
    (" begin open :x for select tblname,tblowner from tables; end;");
cstmt.registerOutParameter(1, TimesTenTypes.CURSOR);
cstmt.execute();
cursor = ((TimesTenCallableStatement)cstmt).getCursor(1);

// Use the cursor as you would any other ResultSet object.
while(cursor.next()){
    System.out.println(cursor.getString(1));
}

// Close the cursor, statement, and connection.
cursor.close();
cstmt.close();
conn.close();
...
```

**Standard REF CURSOR:**

```
...
Connection conn = null;
CallableStatement cstmt = null;
ResultSet rs;
...
cstmt = conn.prepareCall
    (" begin open :x for select tblname,tblowner from tables; end;");
cstmt.registerOutParameter(1, Types.REF_CURSOR);
cstmt.execute();
rs = cstmt.getObject(1, ResultSet.class);
while(rs.next()){
    System.out.println(rs.getString(1));
}

// Close the result set, statement, and connection.
rs.close();
cstmt.close();
conn.close();
...
```

**Note**

If you are evaluating the callable statement with different parameter values in a loop, close the cursor each time at the end of the loop. The typical use case is to prepare the statement, then, in the loop, set parameters, execute the statement, process the cursor, and close the cursor.

## Working with DML Returning (RETURNING INTO Clause)

You can use a `RETURNING INTO` clause, referred to as *DML returning*, with an `INSERT`, `UPDATE`, or `DELETE` statement to return specified items from a row that was affected by the action. This eliminates the need for a subsequent `SELECT` statement and separate round trip, in case, for example, you want to confirm what was affected by the action.

With TimesTen, DML returning is limited to returning items from a single-row operation. The clause returns the items into a list of output parameters.

`TimesTenPreparedStatement`, an extension of the standard `PreparedStatement` interface, supports DML returning. Use the `TimesTenPreparedStatement` method `registerReturnParameter()` to register the return parameters.

```
void registerReturnParameter(int paramIndex, int sqlType)
```

As with the `registerOutParameter()` method discussed in [Working with Output and Input/Output Parameters](#), this method has a signature that enables you to optionally specify a maximum size for `CHAR`, `VARCHAR`, `NCHAR`, `NVARCHAR`, `BINARY`, or `VARBINARY` data. This avoids possible inefficiency where TimesTen would otherwise allocate memory to hold the largest possible value. For `CHAR`, `VARCHAR`, `NCHAR`, and `NVARCHAR`, the unit of size is number of characters. For `BINARY` and `VARBINARY`, it is bytes.

```
void registerReturnParameter(int paramIndex, int sqlType, int maxSize)
```

Use the `TimesTenPreparedStatement` method `getReturnResultSet()` to retrieve the return parameters, returning a `ResultSet` instance.

Be aware of the following restrictions when using `RETURNING INTO` in TimesTen JDBC.

- The `getReturnResultSet()` method must not be invoked more than once. Otherwise, the behavior is indeterminate.
- `ResultSetMetaData` is not supported for the result set returned by `getReturnResultSet()`.
- Streaming methods such as `getCharacterStream()` are not supported for the result set returned by `getReturnResultSet()`.
- There is no batch support for DML returning.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference* for additional information about the TimesTen JDBC classes, interfaces, and methods discussed here.

SQL syntax and restrictions for the `RETURNING INTO` clause in TimesTen are documented as part of the `INSERT`, `UPDATE`, and `DELETE` documentation in *Oracle TimesTen In-Memory Database SQL Reference*.

Refer to RETURNING INTO Clause in *Oracle Database PL/SQL Language Reference* for general information about DML returning.

✓ **Tip**

Check for SQL warnings after executing the TimesTen prepared statement. In the event of a warning, output parameters are undefined. See [Error Handling](#).

This example shows how to use DML returning with a `TimesTenPreparedStatement` instance, returning the name and age for a row that is inserted.

```
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.SQLWarning;
import java.sql.Types;
import com.timesten.jdbc.TimesTenPreparedStatement;

Connection conn = null;
...

// Insert into a table and return results
TimesTenPreparedStatement pstmt =
    (TimesTenPreparedStatement)conn.prepareStatement
    ("insert into tabl values(?,?) returning name, age into ?,?");

// Populate table
pstmt.setString(1,"John Doe");
pstmt.setInt(2, 65);

/* register returned parameter
 * in this case the maximum size of name is 100 chars
 */
pstmt.registerReturnParameter(3, Types.VARCHAR, 100);
pstmt.registerReturnParameter(4, Types.INTEGER);

// process the DML returning statement
int count = pstmt.executeUpdate();

/* Check warnings; if there are warnings, values of DML RETURNING INTO
parameters are undefined. */
SQLWarning wn;
boolean warningFlag = false;
if ((wn = pstmt.getWarnings() ) != null) {
    do {
        warningFlag = true;
        System.out.println(wn);
        wn = wn.getNextWarning();
    } while(wn != null);
}

if (!warningFlag) {
    if (count>0)
    {
        ResultSet rset = pstmt.getReturnResultSet(); //rset not null, not
```

```
empty
    while(rset.next())
    {
        String name = rset.getString(1);
        int age = rset.getInt(2);
        System.out.println("Name " + name + " age " + age);
    }
}
```

## Working with Rowids

Each row in a table has a unique identifier known as its *rowid*. An application can retrieve the rowid of a row from the ROWID pseudocolumn.

A rowid value can be represented in either binary or character format, with the binary format taking 12 bytes and the character format 18 bytes.

TimesTen supports the `java.sql.RowId` interface and `Types.ROWID` type.

You can use the following `ResultSet` methods to retrieve a rowid:

- `RowId getRowId(int columnIndex)`
- `RowId getRowId(String columnLabel)`

You can use the following `PreparedStatement` method to set a rowid:

- `setRowId(int parameterIndex, RowId x)`

An application can specify literal rowid values in SQL statements, such as in `WHERE` clauses, as `CHAR` constants enclosed in single quotes.

Refer to ROWID Data Type and ROWID Pseudocolumn in *Oracle TimesTen In-Memory Database SQL Reference* for additional information about rowids and the ROWID data type, including usage and lifecycle.

### Note

TimesTen does not support the PL/SQL type UROWID.

## Large Objects (LOBs)

TimesTen Classic supports LOBs (large objects), specifically CLOBs (character LOBs), NCLOBs (national character LOBs), and BLOBs (binary LOBs).

This section provides a brief overview of LOBs and discusses their use in JDBC, covering the following topics:

- [About LOBs](#)
- [LOB Objects in JDBC](#)
- [Differences Between TimesTen LOBs and Oracle Database LOBs](#)
- [LOB Factory Methods](#)
- [LOB Getter and Setter Methods](#)

- [TimesTen LOB Interface Methods](#)
- [LOB Prefetching](#)
- [Passthrough LOBs](#)

You can also refer to the following.

- LOB Data Types in *Oracle TimesTen In-Memory Database SQL Reference*
- *Oracle Database SecureFiles and Large Objects Developer's Guide* for general information about programming with LOBs (but not specific to TimesTen functionality)

## About LOBs

A LOB is a large binary object (BLOB) or character object (CLOB or NCLOB). In TimesTen, a BLOB can be up to 16 MB and a CLOB or NCLOB up to 4 MB.

LOBs in TimesTen have essentially the same functionality as in Oracle Database, except as noted otherwise. (See [Differences Between TimesTen LOBs and Oracle Database LOBs](#).)

LOBs may be either persistent or temporary. A persistent LOB exists in a LOB column in the database. A temporary LOB exists only within an application. There are also circumstances where a temporary LOB is created implicitly by TimesTen. For example, if a `SELECT` statement selects a LOB concatenated with an additional string of characters, TimesTen creates a temporary LOB to contain the concatenated data.

## LOB Objects in JDBC

In JDBC, a LOB object—`Blob`, `Clob`, or `NClob` instance—is implemented using a SQL LOB locator (`BLOB`, `CLOB`, or `NCLOB`), which means a LOB object contains a logical pointer to the LOB data rather than the data itself.

### ✓ Tip

- Because LOB objects do not remain valid past the end of the transaction in TimesTen, it is not feasible to use them with autocommit enabled. You would receive errors about LOBs being invalidated.
- LOB manipulations through APIs that use LOB locators result in usage of TimesTen temporary space. Any significant number of such manipulations may necessitate a size increase for the TimesTen temporary data region. See `TempSize` in *Oracle TimesTen In-Memory Database Reference*.

An application can use the JDBC API to instantiate a temporary LOB explicitly, for use within the application, then to free the LOB when done with it. Temporary LOBs are stored in the TimesTen temporary data region.

To update a persistent LOB, your transaction must have an exclusive lock on the row containing the LOB. You can accomplish this by selecting the LOB with a `SELECT ... FOR UPDATE` statement. This results in a writable locator. With a `SELECT` statement, the locator is read-only. Read-only and writable locators behave as follows:

- A read-only locator is *read consistent*, meaning that throughout its lifetime, it sees only the contents of the LOB as of the time it was selected. Note that this would include any uncommitted updates made to the LOB within the same transaction before the LOB was selected.

- A writable locator is updated with the latest data from the database each time a write is made through the locator. So each write is made to the most current data of the LOB, including updates that have been made through other locators.

The following example details behavior for two writable locators for the same LOB.

1. The LOB column contains "XY".
2. Select locator L1 for update.
3. Select locator L2 for update.
4. Write "Z" through L1 at offset 1.
5. Read through locator L1. This would return "ZY".
6. Read through locator L2. This would return "XY", because L2 remains read-consistent until it is used for a write.
7. Write "W" through L2 at offset 2.
8. Read through locator L2. This would return "ZW". Prior to the write in the preceding step, the locator was updated with the latest data ("ZY").

## Differences Between TimesTen LOBs and Oracle Database LOBs

There are significant differences between LOBs in TimesTen and LOBs in Oracle Database.

- In TimesTen, LOB objects do not remain valid past the end of the transaction. All LOB objects are invalidated after a commit or rollback, whether explicit or implicit. This includes after any autocommit (making it infeasible to use LOBs with autocommit enabled), or after any DDL statement.
- TimesTen does not support BFILES, SecureFiles, reads and writes for arrays of LOBs, or callback functions for LOBs.
- TimesTen does not support binding associative arrays of LOBs.
- TimesTen does not support batch processing of LOBs.
- Relevant to BLOBs, there are differences in the usage of hexadecimal literals in TimesTen. See the description of *HexadecimalLiteral* in Constants in *Oracle TimesTen In-Memory Database SQL Reference*.

## LOB Factory Methods

TimesTen supports the standard `Connection` methods `createBlob()`, `createClob()`, and `createNClob()`.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

### ✓ Tip

In TimesTen, creation of a LOB object results in creation of a database transaction if one is not already in progress. You must execute a commit or rollback to close the transaction.

## LOB Getter and Setter Methods

You can access LOBs through getter and setter methods that are defined by the standard `java.sql.ResultSet`, `PreparedStatement`, and `CallableStatement` interfaces, just as they are for other data types.

Use the appropriate `getXXX()` method to retrieve a LOB result or output parameter or `setXXX()` method to bind a LOB input parameter:

- **ResultSet getter methods:** There are `getBlob()` methods, `getClob()` methods, and `getNClob()` methods where you can specify the LOB to retrieve according to either column name or column index.

You can also use `getObject()` to retrieve a `Blob`, `Clob`, or `NClob` object.

- **PreparedStatement setter methods:** There is a `setBlob()` method, `setClob()` method, and `setNClob()` method where you can input the `Blob`, `Clob`, or `NClob` instance and the parameter index to bind an input parameter.

You can also use `setObject()` to bind a `Blob`, `Clob`, or `NClob` input parameter.

There are also `setBlob()` methods where instead of a `Blob` instance, you specify an `InputStream` instance, or an `InputStream` instance and length.

There are `setClob()` and `setNClob()` methods where instead of a `Clob` or `NClob` instance, you specify a `Reader` instance, or a `Reader` instance and length.

- **CallableStatement getter methods:** There are `getBlob()` methods, `getClob()` methods, and `getNClob()` methods where you can retrieve the LOB output parameter according to either parameter name or parameter index.

You can also use `getObject()` to retrieve a `Blob`, `Clob`, or `NClob` output parameter.

You must also register an output parameter from a `CallableStatement` object. The `registerOutParameter()` method takes the parameter index along with the SQL type: `Types.BLOB`, `Types.CLOB`, or `Types.NCLOB`.

- **CallableStatement setter methods:** These are identical to (inherited from) `PreparedStatement` setter methods.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## TimesTen LOB Interface Methods

You can cast a `Blob` instance to `com.timesten.jdbc.TimesTenBlob`, a `Clob` instance to `com.timesten.jdbc.TimesTenClob`, and an `NClob` instance to `com.timesten.jdbc.TimesTenNClob`.

These interfaces support methods specified by the `java.sql.Blob`, `Clob`, and `NClob` interfaces.

The following list summarizes `Blob` features.

- The `isPassthrough()` method, a TimesTen extension, indicates whether the BLOB is a passthrough LOB from Oracle Database.
- Free `Blob` resources when the application is done with it.
- Retrieve the BLOB value as a binary stream. There are methods to retrieve it in whole or in part.
- Retrieve all or part of the BLOB value as a byte array.

- Return the number of bytes in the BLOB.
- Retrieve a stream to be used to write binary data to the BLOB, beginning at the specified position. This overwrites existing data.
- Specify an array of bytes to write to the BLOB, beginning at the specified position, and return the number of bytes written. This overwrites existing data. There are methods to write either all or part of the array.
- Truncate the BLOB to the specified length.

The following list summarizes `Clob` and `NClob` features.

- The `isPassthrough()` method, a TimesTen extension, indicates whether the CLOB or NCLOB is a passthrough LOB from Oracle Database.
- Free `Clob` or `NClob` resources when the application is done with it.
- Retrieve the CLOB or NCLOB as an ASCII stream.
- Retrieve the CLOB or NCLOB as a `java.io.Reader` object (or as a stream of characters). There are methods to retrieve it in whole or in part.
- Retrieve a copy of the specified substring in the CLOB or NCLOB, beginning at the specified position for up to the specified length.
- Return the number of characters in the CLOB or NCLOB.
- Retrieve a stream to be used to write ASCII characters to the CLOB or NCLOB, beginning at the specified position. This overwrites existing data.
- Specify a Java `String` value to write to the CLOB or NCLOB, beginning at the specified position. This overwrites existing data. There are methods to write either all or part of the `String` value.
- Truncate the CLOB or NCLOB to the specified length.

#### Note

- For methods that write data to a LOB, the size of the LOB does not change other than in the circumstance where from the specified position there is less space available in the LOB than there is data to write. In that case, the LOB size increases enough to accommodate the data.
- If the value specified for the position at which to write to a LOB is greater than LOB length + 1, the behavior is undefined.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## LOB Prefetching

To reduce round trips to the server in client/server connections, LOB prefetching is enabled by default when you fetch a LOB from the database. The default prefetch size is 4000 bytes for BLOBs or 4000 characters for CLOBs or NCLOBs.

You can use the `TimesTenConnection` property `CONNECTION_PROPERTY_DEFAULT_LOB_PREFETCH_SIZE` to set a different default value that applies to any statement in the connection. Use a value of -1 to disable LOB prefetching by default for the connection, 0 (zero) to enable LOB prefetching for only metadata by default, or any value

greater than 0 to specify the number of bytes for BLOBs or characters for CLOBs and NCLOBs to be prefetched by default along with the LOB locator during fetch operations.

At the statement level, you can use the following `TimesTenStatement` methods to manipulate the prefetch size and override the default value from the connection:

- `setLobPrefetchSize(int)`: Set a new LOB prefetch value for the statement.
- `int getLobPrefetchSize()`: Return the current LOB prefetch value that applies to the statement (either a value set in the statement itself or the default value from the connection, as applicable).

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## Passthrough LOBs

Passthrough LOBs, which are LOBs in Oracle Database accessed through TimesTen, are exposed as TimesTen LOBs and are supported by TimesTen in much the same way that any TimesTen LOB is supported.

Note the following:

- As noted in [TimesTen LOB Interface Methods](#), the `TimesTenBlob`, `TimesTenClob`, and `TimesTenNClob` interfaces specify the following method to indicate whether the LOB is a passthrough LOB:

```
boolean isPassthrough()
```

- TimesTen LOB size limitations do not apply to storage of LOBs in the Oracle database through passthrough.
- As with TimesTen local LOBs, a passthrough LOB object does not remain valid past the end of the transaction.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## Using CALL to Execute Procedures and Functions

TimesTen supports each of the following syntax formats from any of its programming interfaces to call PL/SQL procedures (*procname*) or PL/SQL functions (*funcname*) that are standalone or part of a package, or to call TimesTen built-in procedures (*procname*).

```
CALL procname[(argumentlist)]
```

```
CALL funcname[(argumentlist)] INTO :returnparam
```

```
CALL funcname[(argumentlist)] INTO ?
```

TimesTen JDBC also supports each of the following syntax formats:

```
{ CALL procname[(argumentlist)] }
```

```
{ ? = [CALL] funcname[(argumentlist)] }
```

```
{ :returnparam = [CALL] funcname[(argumentlist)] }
```

You can execute procedures and functions through the `CallableStatement` interface, with a prepare step first when appropriate (such as when a result set is returned).

The following example calls the TimesTen built-in procedure `ttCkpt`. (Also see the `ttCkpt` example later in this section for a more complete example with JDBC syntax.)

```
CallableStatement.execute("call ttCkpt")
```

The following example calls the TimesTen built-in procedure `ttDataStoreStatus`. A prepare call is used because this procedure produces a result set. (Also see the `ttDataStoreStatus` example later in this section for a more complete example with JDBC syntax.)

```
CallableStatement cStmt = null;
cStmt = conn.prepareCall("call ttDataStoreStatus");
cStmt.execute();
```

The following examples call a PL/SQL procedure `myproc` with two parameters.

```
cStmt.execute("{ call myproc(:param1, :param2) }");
cStmt.execute("{ call myproc(?, ?) }");
```

The following shows several ways to call a PL/SQL function `myfunc`.

```
cStmt.execute("CALL myfunc() INTO :retparam");
cStmt.execute("CALL myfunc() INTO ?");
cStmt.execute("{ :retparam = myfunc() }");
cStmt.execute("{ ? = myfunc() }");
```

See `CALL` in *Oracle TimesTen In-Memory Database SQL Reference*.

### Note

A user's own procedure takes precedence over a TimesTen built-in procedure with the same name, but it is best to avoid such naming conflicts.

This example calls the `ttCkpt` procedure to initiate a fuzzy checkpoint.

```
Connection conn = null;
CallableStatement cStmt = null;
.....
cStmt = conn.prepareCall("{ Call ttCkpt }");
cStmt.execute();
conn.commit();           // commit the transaction
```

Be aware that the `ttCkpt` built-in procedure requires `ADMIN` privilege. Refer to `ttCkpt` in *Oracle TimesTen In-Memory Database Reference*.

The following example calls the `ttDataStoreStatus` procedure and prints out the returned result set. For built-in procedures that return results, you can use the `getXXX()` methods of the `ResultSet` interface to retrieve the data, as shown.

Contrary to the advice given in [Working with TimesTen Result Sets: Hints and Restrictions](#), this example uses a `getString()` call on the `ResultSet` object to retrieve the `Context` field, which is a binary. This is because the output is printed, rather than used for processing. If you do not want to print the `Context` value, you can achieve better performance by using the `getBytes()` method instead.

```
ResultSet rs;

CallableStatement cStmt = conn.prepareCall("{ Call ttDataStoreStatus }");

if (cStmt.execute() == true) {
    rs = cStmt.getResultSet();
    System.out.println("Fetching result set...");
    while (rs.next()) {
        System.out.println("\n Database: " + rs.getString(1));
        System.out.println(" PID: " + rs.getInt(2));
        System.out.println(" Context: " + rs.getString(3));
        System.out.println(" ConType: " + rs.getString(4));
        System.out.println(" memoryID: " + rs.getString(5));
    }
    rs.close();
}
cStmt.close();
```

## Setting a Timeout or Threshold for Executing SQL Statements

TimesTen offers two ways to limit the time for SQL statements to execute, applying to any `execute()`, `executeBatch()`, `executeQuery()`, `executeUpdate()`, or `next()` call. One way is to set a timeout and the other way is to set a threshold.

These mechanisms are described in the following sections:

- [Setting a Timeout Duration for SQL Statements](#)
- [Setting a Threshold Duration for SQL Statements](#)

### Setting a Timeout Duration for SQL Statements

TimesTen allows you to specify a SQL timeout, where if the timeout duration is reached, the statement stops executing and an error is thrown.

You can specify this timeout value for a connection, and therefore any statement on the connection, by using either the `SQLQueryTimeout` general connection attribute (in seconds) or the `SQLQueryTimeoutMsec` general connection attribute (in milliseconds). The default value of each is 0, for no timeout. (Also see `SQLQueryTimeout` and `SQLQueryTimeoutMsec` in *Oracle TimesTen In-Memory Database Reference*.)

Despite the names, these timeout values apply to any executable SQL statement, not just queries.

For a particular statement, you can override the `SQLQueryTimeout` setting by calling the `Statement` method `setQueryTimeout()`.

The query timeout limit has effect only when the SQL statement is actively executing. A timeout does not occur during the commit or rollback phase of an operation. For those transactions that update, insert or delete a large number of rows, the commit or rollback phases may take a long time to complete. During that time the timeout value is ignored.

See Choose SQL and PL/SQL Timeout Values in *Oracle TimesTen In-Memory Database Operations Guide* for considerations regarding the SQL query timeout with respect to other timeout settings.

#### Note

If both a lock timeout value and a SQL query timeout value are specified, the lesser of the two values causes a timeout first. Regarding lock timeouts, you can refer to `ttLockWait` (built-in procedure) or `LockWait` (general connection attribute) in *Oracle TimesTen In-Memory Database Reference*, or to Check for Deadlocks and Timeouts in *Oracle TimesTen In-Memory Database Monitoring and Troubleshooting Guide*.

## Setting a Threshold Duration for SQL Statements

TimesTen allows you to set a SQL threshold time period, in seconds. If the threshold is reached, a warning is written to the support log but execution continues.

Despite the name, this threshold applies to any JDBC call executing a SQL statement, not just queries.

By default, the application obtains the threshold value from the `QueryThreshold` general connection attribute setting, for which the default is 0 (no warnings). You can override the threshold for a JDBC `Connection` object by including the `QueryThreshold` attribute in the connection URL for the database. For example, to set `QueryThreshold` to a value of 5 seconds for the `myDSN` database:

```
jdbc:timesten:direct:dsn=myDSN;QueryThreshold=5
```

You can also use the `setQueryTimeThreshold()` method of a `TimesTenStatement` object to set the threshold. This overrides the connection attribute setting and the `Connection` object setting.

You can retrieve the current threshold value by using the `getQueryTimeThreshold()` method of the `TimesTenStatement` object.

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## Configuring the Result Set Buffer Size in Client/Server Using JDBC

For data returned from a `SELECT` statement in client/server, the buffer size for the data returned to the client is configurable to allow adjustments for better performance. (In earlier releases, the buffer size could not be changed.)

The buffer size can be set in terms of either rows of data or bytes of data. The lower limit takes precedence.

For client/server, the TimesTen implementation of the `setFetchSize()` and `getFetchSize()` methods on `Statement` and `ResultSet` objects allow you to set or get the maximum number of rows of data:

- `void setFetchSize(int rows)`

- `int getFetchSize()`

For portability, this is the preferred way to set the buffer size. The default value is 8192. If you set this to a large value, you may also need to set the maximum number of bytes, as shown below, to ensure that will not be a constraint.

In addition, these methods are in `TimesTenStatement`:

- `void setTtNetMsgMaxRows(int rows)`
- `int getTtNetMsgMaxRows()`
- `void setTtNetMsgMaxBytes(int bytes)`
- `int getTtNetMsgMaxBytes()`

And the set methods are also in `TimesTenConnection`:

- `void setTtNetMsgMaxRows(int rows)`
- `void setTtNetMsgMaxBytes(int bytes)`

The default number of rows is 8192; the default number of bytes is 2097152, or 2 MB. It is suggested to use one limit and set the other to a value high enough to ensure that it is not reached first.

When you use one of the TimesTen setter methods on a connection object, the value will apply to any future statement objects created on the connection and also to any existing statement objects on the connection. It is recommended, though, to set them at statement level (or at connection level only to serve as initial values for statement handles to be created).

#### Note

- These settings correspond to TimesTen connection attributes `TT_NetMsgMaxRows` and `TT_NetMsgMaxBytes`, which you can set in a TimesTen connection string or DSN, to serve as initial values for any statements created on the connection.
- The minimum value of each setting is 1 and at least one row is always returned. Setting either rows or bytes to a value of 0 results in the default value being used for that setting. There are no maximum settings other than the maximum value of the data type (32-bit signed integer).
- If a client version that supports these attributes connects to a server version that does not, any settings are silently ignored.

## Features for Use with Cache

This section discusses features related to the use of cache operations.

- [Setting the Oracle Database Password](#)
- [Setting Temporary Passthrough Level with the `ttOptSetFlag` Built-In Procedure](#)
- [Determining Passthrough Status](#)
- [Managing Cache Groups](#)

## Setting the Oracle Database Password

The `OraclePassword` attribute maps to the Oracle Database password. You can use the `TimesTenDataSource` method `setOraclePassword()` to set the Oracle Database password.

See [Connecting to the Database](#) for an example.

## Setting Temporary Passthrough Level with the `ttOptSetFlag` Built-In Procedure

TimesTen provides the `ttOptSetFlag` built-in procedure for setting various flags, including the `PassThrough` flag to temporarily set the passthrough level.

You can use `ttOptSetFlag` to set `PassThrough` in a JDBC application as in the following sample statement, which sets the passthrough level to 1. The setting affects all statements that are prepared until the end of the transaction.

```
pstmt = conn.prepareStatement("call ttOptSetFlag('PassThrough', 1)");
```

The example that follows has samples of code that accomplish these steps:

1. Create a prepared statement (a `PreparedStatement` instance `thePassThroughStatement`) that calls `ttOptSetFlag` using a bind parameter for passthrough level.
2. Define a method `setPassthrough()` that takes a specified passthrough setting, binds it to the prepared statement, then executes the prepared statement to call `ttOptSetFlag` to set the passthrough level.

```
thePassThroughStatement =  
    theConnection.prepareStatement("call  
ttOptSetFlag('PassThrough', ?)");  
...  
private void setPassthrough(int level) throws SQLException{  
    thePassThroughStatement.setInt(1, level);  
    thePassThroughStatement.execute();  
}
```

See `ttOptSetFlag` and `PassThrough` in *Oracle TimesTen In-Memory Database Reference*. See *Setting a Passthrough Level in Oracle TimesTen In-Memory Database Cache Guide* for information about passthrough settings.

## Determining Passthrough Status

You can call the `TimesTenPreparedStatement` method `getPassThroughType()` to determine whether a SQL statement is to be executed in the TimesTen database or passed through to the Oracle database for execution.

```
PassThroughType getPassThroughType()
```

The return type, `TimesTenPreparedStatement.PassThroughType`, is an enumeration type for values of the TimesTen `PassThrough` connection attribute.

You can make this call after preparing the SQL statement. It is useful with `PassThrough` settings of 1 or 2, where the determination of whether a statement is actually passed through is not made until compilation time.

See *Setting a Passthrough Level in Oracle TimesTen In-Memory Database Cache Guide*.

## Managing Cache Groups

In TimesTen, following the execution of a `FLUSH CACHE GROUP`, `LOAD CACHE GROUP`, `REFRESH CACHE GROUP`, or `UNLOAD CACHE GROUP` statement, the `Statement` method `getUpdateCount()` returns the number of cache instances that were flushed, loaded, refreshed, or unloaded.

See *Determining the Number of Cache Instances Affected by an Operation in Oracle TimesTen In-Memory Database Cache Guide*.

## Features for Use with Replication

For TimesTen Classic applications that employ replication, you can improve performance by using *parallel replication*, which uses multiple threads acting in parallel to replicate and apply transactional changes to databases in a replication scheme.

TimesTen supports the following types of parallel replication:

- Automatic parallel replication (`ReplicationApplyOrdering=0`): Parallel replication over multiple threads that automatically enforces transactional dependencies and all changes applied in commit order. This is the default.
- Automatic parallel replication with disabled commit dependencies (`ReplicationApplyOrdering=2`): Parallel replication over multiple threads that automatically enforces transactional dependencies, but does not enforce transactions to be committed in the same order on the subscriber database as on the master database. In this mode, you can optionally specify replication tracks.

See *Configuring Parallel Replication in Oracle TimesTen In-Memory Database Replication Guide*.

For JDBC applications that use parallel replication and specify replication tracks, you can specify the track number for transactions on a connection through the following `TimesTenConnection` method. (Alternatively, use the general connection attribute `ReplicationTrack` or the `ALTER SESSION` parameter `REPLICATION_TRACK`.)

- `void setReplicationTrack(int track)`

`TimesTenConnection` also has the corresponding getter method:

- `int getReplicationTrack()`

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

## Error Handling

There are methods to check for, identify, and handle errors in a TimesTen Java application.

For a list of the errors that TimesTen returns and what to do if errors are encountered, see the [Error Help portal](#).

This section includes the following topics.

- [Error and Warning Levels](#)
- [Reporting Errors and Warnings](#)
- [Catching and Responding to Specific Errors](#)
- [Rolling Back Failed Transactions](#)

- [Retrying After Transient Errors \(JDBC\)](#)

## Error and Warning Levels

When operations are not completely successful, TimesTen can return a fatal error, a non-fatal error, or a warning.

The sections that follow discuss error or warning situations and how to handle them:

- [Fatal Errors](#)
- [Non-Fatal Errors](#)
- [Warnings](#)
- [Abnormal Termination](#)

## Fatal Errors

Fatal errors make the database inaccessible until it can be recovered. When a fatal error occurs, all database connections are required to disconnect. No further operations may complete.

Fatal errors are indicated by TimesTen error codes 846 and 994. Error handling for these errors should be different from standard error handling. In particular, the code should roll back the current transaction and, to avoid out-of-memory conditions in the server, disconnect from the database. Shared memory from the old TimesTen instance is not freed until all connections that were active at the time of the error have disconnected. Inactive applications still connected to the old TimesTen instance may have to be manually terminated.

When fatal errors occur, TimesTen performs the full cleanup and recovery procedure:

- Every connection to the database is invalidated, a new memory segment is allocated and applications are required to disconnect.
- The database is recovered from the checkpoint and transaction log files upon the first subsequent initial connection.
  - The recovered database reflects the state of all durably committed transactions and possibly some transactions that were committed non-durably.
  - No uncommitted or rolled back transactions are reflected.

## Non-Fatal Errors

Non-fatal errors include errors such as an `INSERT` statement that violates unique constraints. This category also includes some classes of application and process failures.

TimesTen returns non-fatal errors through the typical error-handling process. Application should check for errors and appropriately handle them.

When a database is affected by a non-fatal error, an error may be returned and the application should take appropriate action.

An application can handle non-fatal errors by modifying its actions or, in some cases, by rolling back one or more offending transactions, as described in [Rolling Back Failed Transactions](#).

Also see [Reporting Errors and Warnings](#), which follows shortly.

**Note**

If a `ResultSet`, `Statement`, `PreparedStatement`, `CallableStatement` or `Connection` operation results in a database error, it is a good practice to call the `close()` method for that object.

## Warnings

TimesTen returns a warning when something occurs that is unexpected but not considered very serious.

Here are some examples of events that cause TimesTen to issue a warning:

- A checkpoint failure
- Use of a deprecated TimesTen feature
- Truncation of some data
- Execution of a recovery process upon connect
- Replication return receipt timeout

You should always have code that checks for warnings, as they can indicate application problems.

Also see [Reporting Errors and Warnings](#) immediately below.

## Abnormal Termination

In some cases, such as with a process failure, an error cannot be returned, so TimesTen automatically rolls back the transactions of the failed process.

## Reporting Errors and Warnings

Your application should check for and report all errors and warnings that can be returned on every call. This saves considerable time and effort during development and debugging.

A `SQLException` object is generated if there are one or more database access errors and a `SQLWarning` object is generated if there are one or more warning messages. A single call may return multiple errors or warnings or both, so your application should report all errors or warnings in the returned `SQLException` or `SQLWarning` objects.

Multiple errors or warnings are returned in linked chains of `SQLException` or `SQLWarning` objects. The examples below demonstrate how you might iterate through the lists of returned `SQLException` and `SQLWarning` objects to report all of the errors and warnings, respectively.

The following method prints out the content of all exceptions in the linked `SQLException` objects.

```
static int reportSQLExceptions(SQLException ex)
{
    int errCount = 0;
    if (ex != null) {
        errStream.println("\n--- SQLException caught ---");
        ex.printStackTrace();
    }
}
```

```

        while (ex != null) {
            errStream.println("SQL State: " + ex.getSQLState());
            errStream.println("Message: " + ex.getMessage());
            errStream.println("Error Code: " + ex.getErrorCode());
            errCount++;
            ex = ex.getNextException();
            errStream.println();
        }
    }

    return errCount;
}

```

This next method prints out the content of all warning in the linked `SQLWarning` objects.

```

static int reportSQLWarnings(SQLWarning wn)
{
    int warnCount = 0;

    while (wn != null) {
        errStream.println("\n--- SQL Warning ---");
        errStream.println("SQL State: " + wn.getSQLState());
        errStream.println("Message: " + wn.getMessage());
        errStream.println("Error Code: " + wn.getErrorCode());

        // is this a SQLWarning object or a DataTruncation object?
        if (wn instanceof DataTruncation) {
            DataTruncation trn = (DataTruncation) wn;
            errStream.println("Truncation error in column: " +
                trn.getIndex());
        }
        warnCount++;
        wn = wn.getNextWarning();
        errStream.println();
    }
    return warnCount;
}

```

## Catching and Responding to Specific Errors

In some situations it may be desirable for your application to respond in a certain way to a specific SQL state or TimesTen error code. You can use the `SQLException` method `getSQLState()` to return the SQL state and the `getErrorCode()` method to return TimesTen error codes.

Also refer to the entry for `TimesTenVendorCode` in *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference*.

The TimesTen Classic s Quick Start ample applications require you to load the schema before executing the samples. The following `catch` statement alerts the user that `appuser` has not

been loaded or has not been refreshed by detecting ODBC error S0002 and TimesTen error 907:

```
catch (SQLException ex) {
    if (ex.getSQLState().equalsIgnoreCase("S0002")) {
        errStream.println("\nError: The table appuser.customer " +
            "does not exist.\n\t Please reinitialize the database.");
    } else if (ex.getErrorCode() == 907) {
        errStream.println("\nError: Attempting to insert a row " +
            "with a duplicate primary key.\n\tPlease reinitialize the database.");
    }
}
```

You can use the `TimesTenVendorCode` interface to detect the errors by their name, rather than their number.

Consider this example:

```
ex.getErrorCode() == com.timesten.jdbc.TimesTenVendorCode.TT_ERR_KEYEXISTS
```

The following is equivalent:

```
ex.getErrorCode() == 907
```

## Rolling Back Failed Transactions

In some situations, such as recovering from a deadlock or lock timeout, you should explicitly roll back the transaction using the `Connection` method `rollback()`.

```
try {
    if (conn != null && !conn.isClosed()) {
        // Rollback any transactions in case of errors
        if (retcode != 0) {
            try {
                System.out.println("\nEncountered error. Rolling back transaction");
                conn.rollback();
            } catch (SQLException ex) {
                reportSQLExceptions(ex);
            }
        }
    }

    System.out.println("\nClosing the connection\n");
    conn.close();
} catch (SQLException ex) {

    reportSQLExceptions(ex);
}
```

A transaction rollback consumes resources and the entire transaction is in effect wasted. To avoid unnecessary rollbacks, design your application to avoid contention and check the application or input data for potential errors before submitting it.

**Note**

If your application terminates, crashes, or disconnects in the middle of an active transaction, TimesTen automatically rolls back the transaction.

## Retrying After Transient Errors (JDBC)

TimesTen automatically resolves most transient errors, but if your application detects the following `SQLSTATE` value, it is suggested to retry the current transaction.

- `TT005`: Transient transaction failure due to unavailability of resource. Roll back the transaction and try it again.

**Note**

Search the entire error stack for errors returning these error types before deciding whether it is appropriate to retry.

This is returned by the `getSQLState()` method of the `SQLException` class and may be encountered by method calls from any of the following JDBC types:

- `Connection`
- `Statement`
- `PreparedStatement`
- `CallableStatement`
- `ResultSet`
- `Connection`
- `Statement`
- `PreparedStatement`
- `CallableStatement`
- `ResultSet`

Here is an example:

```
// Database connection object
Connection      dbConn;

// Open the connection to the database
...

// Disable auto-commit
dbConn.setAutoCommit( false );

...

// Prepre the SQL statements
PreparedStatement stmtQuery = dbConn.prepareStatement("SELECT ...");
```

```
PreparedStatement stmtUpdate = dbConn.prepare("UPDATE ...");

...

// Set max retries for transaction to 5
int retriesLeft = 5;

// Records outcome
boolean success = false;

// Excute transaction with retries until success or retries exhausted
while ( retriesLeft > 0 )
{
    try {

        // First execute the query

        // Set input values
        stmtQuery.setInt(1, ...);
        stmtQuery.setString(2, ...);

        // Execute and process results
        ResultSet rs = stmtQuery.executeQuery();
        while ( rs.next() )
        {
            int val1 = rs.getInt(1);
            String val2 = rs.getString(2);
            ...
        }
        rs.close();
        rs = null;

        // Now excute the update

        // Set input values
        stmtUpdate.setInt(1,...);
        stmtUpdate.setString(2,...);

        // Execute and check number of rows affected
        int updCount = stmtUpdate.executeUpdate();
        if ( updCount < 1 )
        {
            ...
        }

        // And finally commit
        dbConn.commit();

        // We are done
        success = true;
        break;

    } catch ( SQLException sqe ) {

        if ( sqe.getSQLState().equals("TT005") ) // grid transient error
        {
```

```
        // decrement retry count
        retriesLeft--;
        // and rollback the transaction ready for retry
        try {
            dbConn.rollback();
        } catch ( SQLException sqer ) {
            // This is a fatal error so handle accordingly
        }
    }
    else
    {
        // handle other kinds of error
        ...
    }
} // end of retry loop

if ( ! success )
{
    // Handle the failure
    ...
}
```

### Note

The example in [Failover Delay and Retry Settings](#) also shows how to retry for transient errors.

## JDBC Support for Automatic Client Failover

There are TimesTen JDBC extensions related to automatic client failover.

- [About Automatic Client Failover](#)
- [Features and Functionality of JDBC Support for Automatic Client Failover](#)
- [Configuration of Automatic Client Failover](#)
- [Synchronous Detection of Automatic Client Failover in JDBC](#)
- [Asynchronous Detection of Automatic Client Failover in JDBC](#)
- [JDBC Application Action in the Event of Failover](#)

See Using Automatic Client Failover in *Oracle TimesTen In-Memory Database Operations Guide*. For related information for developers, see ODBC Support for Automatic Client Failover in *Oracle TimesTen In-Memory Database C Developer's Guide*.

## About Automatic Client Failover

Automatic client failover is for use in high availability scenarios. There are two scenarios for TimesTen Classic, one with active standby pair replication and one referred to as *generic automatic client failover*.

If there is a failure of the database or database element to which the client is connected, then failover (connection transfer) to an alternate database or database element occurs:

- For TimesTen with active standby replication, failover is to the new active (original standby) database.
- For TimesTen using generic automatic client failover, where you can ensure that the schema and data are consistent on both databases, failover is to a database from a list that is configured in the client `odbc.ini` file.

A typical use case for generic automatic failover is a set of databases using read-only caching, where each database has the same set of cached data. For example, if you have several read-only cache groups, then you would create the same read-only cache groups on all TimesTen Classic databases included in the list of failover servers. When the client connection fails over to an alternate TimesTen database, the cached data is consistent because TimesTen automatically refreshes the cache data (as needed) from the Oracle database.

Applications are automatically reconnected to the new data database or database element. TimesTen provides features that enable applications to be alerted when this happens, so they can take any appropriate action.

Any of the following error conditions indicates automatic client failover.

- Native error 30105 with SQL state 08006
- Native error 47137

#### Note

- Automatic client failover applies only to client/server connections. The functionality described here does not apply to a direct connection.
- Automatic client failover is complementary to Oracle Clusterware in situations where Oracle Clusterware is used, though the two features are not dependent on each other. You can also refer to *Using Oracle Clusterware to Manage Active Standby Pairs* in *Oracle TimesTen In-Memory Database Replication Guide*.

## Features and Functionality of JDBC Support for Automatic Client Failover

There are TimesTen JDBC features related to client failover, including functionality relating specifically to pooled connections.

- [General Client Failover Features](#)
- [Client Failover Features for Pooled Connections](#)

Refer to *Oracle TimesTen In-Memory Database JDBC Extensions Java API Reference* for additional information about the TimesTen JDBC classes, interfaces, and methods discussed here.

## General Client Failover Features

TimesTen JDBC support for automatic client failover provides two mechanisms for detecting a failover.

- *Synchronous detection*, through a SQL exception: After an automatic client failover, JDBC objects created on the failed connection—such as statements, prepared statements, callable statements, and result sets—can no longer be used. A Java SQL exception is thrown if an application attempts to access any such object. By examining the SQL state

and error code of the exception, you can determine whether the exception is the result of a failover situation.

- *Asynchronous detection*, through an event listener: An application can register a user-defined client failover event listener, which is notified of each event that occurs during the process of a failover.

TimesTen JDBC provides the following features, in package `com.timesten.jdbc`, to support automatic client failover.

- `ClientFailoverEvent` class

This class is used to represent events that occur during a client failover: begin, end, abort, or retry.

- `ClientFailoverEventListener` interface

An application interested in client failover events must have a class that implements this interface, which is the mechanism to listen for client failover events. At runtime, the application must register `ClientFailoverEventListener` instances through the TimesTen connection (see immediately below).

You can use a listener to proactively react to failure detection, such as by refreshing connection pool statement caches, for example.

- Methods in the `TimesTenConnection` interface

This interface specifies the methods `addConnectionEventListener()` and `removeConnectionEventListener()` to register or remove, respectively, a client failover event listener.

- A constant, `TT_ERR_FAILOVERINVALIDATION`, in the `TimesTenVendorCode` interface

This enables you to identify an event as a failover event.

## Client Failover Features for Pooled Connections

TimesTen recommends that applications using pooled connections (`javax.sql.PooledConnection`) or connection pool data sources (`javax.sql.ConnectionPoolDataSource`) use the synchronous mechanism noted previously to handle stale objects on the failed connection.

Java EE application servers manage pooled connections, so applications are not able to listen for events on pooled connections. And application servers do not implement and register an instance of `ClientFailoverEventListener`, because this is a TimesTen extension.

## Configuration of Automatic Client Failover

Refer to *Configuring Automatic Client Failover for TimesTen* in *Oracle TimesTen In-Memory Database Operations Guide*. In TimesTen, failover DSNs must be specifically configured through `TTC_Server2` and `TTC_Servern` connection attributes.

### Note

Setting any of `TTC_Server2`, `TTC_Server_DSN2`, `TTC_Servern`, or `TCP_Port2` implies that you intend to use automatic client failover. For the active standby pair scenario, it also means a new thread is created for your application to support the failover mechanism.

Be aware of these TimesTen connection attributes:

- `TTC_NoReconnectOnFailover`: If this is set to 1 (enabled), TimesTen is instructed to do all the usual client failover processing except for the automatic reconnect. (For example, statement and connection handles are marked as invalid.) This is useful if the application does its own connection pooling or manages its own reconnection to the database after failover. The default value is 0 (reconnect). Also see `TTC_NoReconnectOnFailover` in *Oracle TimesTen In-Memory Database Reference*.
- `TTC_REDIRECT`: If this is set to 0 and the initial connection attempt to the desired database or database element fails, then an error is returned and there are no further connection attempts. This does not affect subsequent failovers on that connection. Also see [#unique\\_104](#) in *Oracle TimesTen In-Memory Database Reference*.
- `TTC_Random_Selection`: For TimesTen Classic using generic automatic client failover, the default setting of 1 (enabled) specifies that when failover occurs, the client randomly selects an alternative server from the list provided in `TTC_Servern` attribute settings. If the client cannot connect to the selected server, it keeps redirecting until it successfully connects to one of the listed servers. With a setting of 0, TimesTen goes through the list of `TTC_Servern` servers sequentially. Also see `TTC_Random_Selection` in *Oracle TimesTen In-Memory Database Reference*.

### Note

If you set any of these in `odbc.ini` or the connection string, the settings are applied to the failover connection. They cannot be set in your application (including by `ALTER SESSION`).

## Synchronous Detection of Automatic Client Failover in JDBC

If, in a failover situation, an application attempts to use objects created on the failed connection, then JDBC throws a SQL exception. The vendor-specific exception code is set to `TimesTenVendorCode.TT_ERR_FAILOVERINVALIDATION`.

Detecting a failover through this mechanism is referred to as synchronous detection. The following example demonstrates this.

```
try {
    // ...
    // Execute a query on a previously prepared statement.
    ResultSet theResultSet = theStatement.executeQuery("select * from dual");
    // ...
} catch (SQLException sqllex) {
    sqllex.printStackTrace();
    if (sqllex.getErrorCode() ==
TimesTenVendorCode.TT_ERR_FAILOVERINVALIDATION) {
        // Automatic client failover has taken place; discontinue use of this
        object.
    }
}
```

## Asynchronous Detection of Automatic Client Failover in JDBC

Asynchronous failover detection requires an application to implement a client failover event listener and register an instance of it on the TimesTen connection.

This section describes the steps involved:

1. [Implement a Client Failover Event Listener.](#)
2. [Register the Client Failover Listener Instance.](#)
3. [Remove the Client Failover Listener Instance.](#)

### Implement a Client Failover Event Listener

TimesTen JDBC provides the `com.timesten.jdbc.ClientFailoverEventListener` interface for use in listening for events.

This interface is highlighted by the following method:

- `void notify(ClientFailoverEvent event)`

To use asynchronous failover detection, you must create a class that implements this interface, then register an instance of the class at runtime on the TimesTen connection (discussed shortly).

When a failover event occurs, TimesTen calls the `notify()` method of the listener instance you registered, providing a `ClientFailoverEvent` instance that you can then examine for information about the event.

The following example shows the basic form of a `ClientFailoverEventListener` implementation.

```
private class MyCFListener implements ClientFailoverEventListener {
    /* Applications can build state system to track states during failover.
       You may want to add methods that talks about readiness of this
Connection
       for processing.
    */
    public void notify(ClientFailoverEvent event) {

        /* Process connection failover type */
        switch(event.getTheFailoverType()) {
            case TT_FO_CONNECTION:
                /* Process session fail over */
                System.out.println("This should be a connection failover type " +
                    event.getTheFailoverType());

                break;

            default:
                break;
        }
        /* Process connection failover events */
        switch(event.getTheFailoverEvent()) {
            case BEGIN:
                System.out.println("This should be a BEGIN event " +
                    event.getTheFailoverEvent());
                /* Applications cannot use Statement, PreparedStatement,
```

```

ResultSet,
    etc. created on the failed Connection any longer.
    */
    break;

case END:
    System.out.println("This should be an END event " +
        event.getTheFailoverEvent());

    /* Applications may want to re-create Statement and
    PreparedStatement
    objects at this point as needed.
    */
    break;

case ABORT:
    System.out.println("This should be an ABORT event " +
        event.getTheFailoverEvent());

    break;

case ERROR:
    System.out.println("This should be an ERROR event " +
        event.getTheFailoverEvent());

    break;

default:
    break;
    }
}
}

```

The `event.getTheFailoverType()` call returns an instance of the nested class `ClientFailoverEvent.FailoverType`, which is an enumeration type. In TimesTen, the only supported value is `TT_FO_CONNECTION`, indicating a connection failover.

The `event.getTheFailoverEvent()` call returns an instance of the nested class `ClientFailoverEvent.FailoverEvent`, which is an enumeration type where the value can be one of the following:

- `BEGIN`, if the client failover has begun
- `END`, if the client failover has completed successfully
- `ERROR`, if the client failover failed but will be retried
- `ABORT`, if the client failover has aborted

## Register the Client Failover Listener Instance

At runtime you must register an instance of your failover event listener class with the TimesTen connection object, so that TimesTen can call the `notify()` method of the listener class as needed for failover events.

`TimesTenConnection` provides the following method for this.

- `void addConnectionEventListener (ClientFailoverEventListener listener)`

Create an instance of your listener class, then register it using this method. The following example establishes the connection and registers the listener. Assume `theDsn` is the JDBC URL for a TimesTen client/server database and `theCFListener` is an instance of your failover event listener class.

```
try {
    /* Assume this is a client/server conn; register for conn failover.
*/
    Class.forName("com.timesten.jdbc.TimesTenClientDriver");
    String url = "jdbc:timesten:client:" + theDsn;
    theConnection = (TimesTenConnection)DriverManager.getConnection(url);
    theConnection.addConnectionEventListener(theCFListener);
    /* Additional logic goes here; connection failover listener is
       called if there is a fail over.
    */
}
catch (ClassNotFoundException cnfex) {
    cnfex.printStackTrace();
}
catch (SQLException sqlex) {
    sqlex.printStackTrace();
}
```

## Remove the Client Failover Listener Instance

The `TimesTenConnection` interface defines the following method to deregister a failover event listener.

- `void removeConnectionEventListener (ClientFailoverEventListener listener)`

Use this method to deregister a listener instance.

## JDBC Application Action in the Event of Failover

There are actions for an application to take if there is a failover.

- [Application Steps for Failover](#)
- [Failover Delay and Retry Settings](#)

## Application Steps for Failover

If you receive any of the error conditions noted at the beginning of automatic client failover discussion in [JDBC Support for Automatic Client Failover](#) in response to an operation in your application, then application failover is in progress. There are actions you can perform to recover after failover.

1. Roll back all transactions on the connection.
2. Clean up all objects from the previous connection. None of the state or objects associated with the previous connection are preserved.
3. Assuming `TTC_NoReconnectOnFailover=0` (the default), sleep briefly, as discussed in the next section, [Failover Delay and Retry Settings](#). If `TTC_NoReconnectOnFailover=1`, then you must instead manually reconnect the application to an alternate database or database element.
4. Recreate and reprepare all objects related to your connection.

5. Restart any in-progress transactions from the beginning.

## Failover Delay and Retry Settings

The reconnection to another database or database element during automatic client failover may take some time. If your application attempts recovery actions before TimesTen has completed its client failover process, you may receive another failover error condition.

See [JDBC Support for Automatic Client Failover](#).

Therefore, your application should place all recovery actions within a loop with a short delay before each subsequent attempt, where the total number of attempts is limited. If you do not limit the number of attempts, the application may appear to freeze if the client failover process does not complete successfully. For example, your recovery loop could use a retry delay of 100 milliseconds with a maximum number of retries limited to 100 attempts. The ideal values depend on your particular application and configuration.

This code example, using the synchronous detection method, illustrates how you might handle the retrying of connection failover errors in a JDBC application. Code not directly relevant is omitted (...).

```
// Database connection object
Connection      dbConn;

// Open the connection to the database
...

// Disable auto-commit
dbConn.setAutoCommit( false );

...

// Prepre the SQL statements
PreparedStatement stmtQuery = dbConn.prepareStatement("SELECT ...");
PreparedStatement stmtUpdate = dbConn.prepareStatement("UPDATE ...");

...

// Set max retries to 100
int retriesLeft = 100;
// and retry delay to 100 ms
int retryDelay = 100;

// Records outcome
boolean success = false;
Boolean needReprepare = false;

// Execute transaction with retries until success or retries exhausted
while ( retriesLeft > 0 )
{
    try {

        // Do we need to re-prepare
        if ( needReprepare )
        {
            Thread.sleep( retryDelay ); // delay before proceeding
            stmtQuery = dbConn.prepareStatement("SELECT ...");
        }
    }
}
```

```
        stmtUpdate = dbConn.prepareStatement("UPDATE ...");
        needReprepare = false;
    }

    // First execute the query

    // Set input values
    stmtQuery.setInt(1, ...);
    stmtQuery.setString(2, ...);

    // Execute and process results
    ResultSet rs = stmtQuery.executeQuery();
    while ( rs.next() )
    {
        int val1 = rs.getInt(1);
        String val2 = rs.getString(2);
        ...
    }
    rs.close();
    rs = null;

    // Now execute the update

    // Set input values
    stmtUpdate.setInt(1,...);
    stmtUpdate.setString(2,...);

    // Execute and check number of rows affected
    int updCount = stmtUpdate.executeUpdate();
    if ( updCount < 1 )
    {
        ...
    }

    // And finally commit
    dbConn.commit();

    // We are done
    success = true;
    break;

} catch ( SQLException sqe ) {

    if ( (sqe.getErrorCode() == 47137) ||
        ( (sqe.getErrorCode() == 30105) &&
          (sqe.getSQLState().equals("08006")) ) )
        // connection failover error
        {
            // decrement retry count
            retriesLeft--;
            // rollback the transaction ready for retry
            dbConn.rollback();
            // and indicate that we need to re-prepare
            needReprepare = true;
        }
    else
```

---

```
        {
            // handle other kinds of error
            ...
        }
    }
} // end of retry loop if ( ! success ){ // Handle the failure ...}
```

# 3

## Using JMS/XLA for Event Management

This chapter documents the TimesTen JMS/XLA API, which you can use in TimesTen Classic to monitor TimesTen for changes to specified tables. JMS/XLA implements Java Message Service (JMS) interfaces to make the functionality of the TimesTen Transaction Log API (XLA) available to Java applications.

This chapter includes the following topics:

- [JMS/XLA Concepts and Features](#)
- [JMS/XLA Usage and Functionality](#)
- [JMS/XLA as a Replication Mechanism](#)

### JMS/XLA Concepts and Features

This section provides an overview of JMS/XLA concepts and features.

- [JMS/XLA Concepts](#)
- [How XLA Reads Records from the Transaction Log](#)
- [XLA and Materialized Views](#)
- [XLA Bookmarks](#)
- [JMS/XLA Configuration File and Topics](#)
- [XLA Updates](#)
- [XLA Update Acknowledgments](#)
- [XLA System Privilege](#)
- [XLA Limitations](#)
- [JMS/XLA and Oracle GDK Dependency](#)

Also, JMS information and resources are available at this location:

<http://www.oracle.com/technetwork/java/jms/index.html>

In addition, the standard JMS API documentation is included with the TimesTen installation at the following location:

`installation_dir/3rdparty/jms1.1/doc/api/index.html`

For information about tuning TimesTen JMS/XLA applications for improved performance, see [Tuning JMS/XLA Applications](#).

### JMS/XLA Concepts

You can use the TimesTen JMS/XLA API, supported by TimesTen Classic, to monitor TimesTen for changes to specified tables in a local database and to receive real-time notification of these changes. The primary purpose of JMS/XLA is as a high-performance, asynchronous alternative to triggers.

TimesTen JMS/XLA support includes two JMS providers: one based on the Jakarta Messaging API version 1.1 and another based Java Message Service API version 1.1. The Jakarta Messaging API version 1.1 is functionally equivalent to the Java Message Service API version 1.1 except for using the namespace `jakarta.jms` instead of `javax.jms`.

JMS/XLA implements JMS interfaces to make the functionality of the TimesTen Transaction Log API (XLA) available to Java applications. JMS/XLA API is a wrapper for XLA, which obtains update records directly from the transaction log buffer or transaction log files, so the records are available until they are read. XLA also allows multiple readers to access transaction log updates simultaneously. See XLA and TimesTen Event Management in *Oracle TimesTen In-Memory Database C Developer's Guide*.

JMS/XLA uses the JMS publish-subscribe interface to provide access to XLA updates. Subscribe to updates by establishing a JMS `Session` instance that provides a connection to XLA and then creating a durable subscriber (`TopicSubscriber`). You can receive and process messages synchronously through the subscriber, or you can implement a listener (`MessageListener`) to process the updates asynchronously.

JMS/XLA is designed for applications that want to monitor a local database. TimesTen and the application receiving the notifications must reside on the same system.

#### Note

In the unlikely event that the TimesTen replication solutions described in Overview of TimesTen Replication in the *Oracle TimesTen In-Memory Database Replication Guide* do not meet your needs, it is possible to use JMS/XLA to build a custom data replication solution.

## How XLA Reads Records from the Transaction Log

As applications modify a database, TimesTen generates transaction log records that describe the changes made to the data and other events such as transaction commits.

New transaction log records are always written to the end of the transaction log buffer as they are generated. Transaction log records are periodically flushed in batches from the log buffer in memory to transaction log files on the file system.

Applications can use XLA to monitor the transaction log for changes to the database. XLA reads through the transaction log, filters the log records, and delivers XLA applications with a list of transaction records that contain the changes to the tables and columns of interest.

XLA sorts the records into discrete transactions. If multiple applications are updating the database simultaneously, transaction log records from the different applications are interleaved in the transaction log.

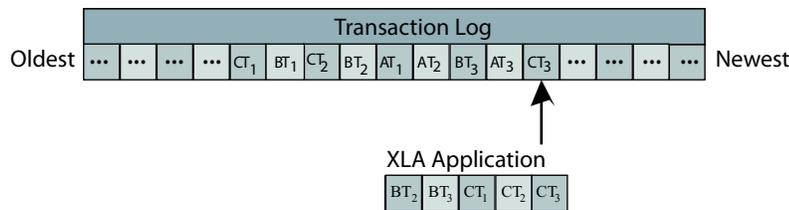
XLA transparently extracts all transaction log records associated with a particular transaction and delivers them in a contiguous list to the application.

Only the records for committed transactions are returned. They are returned in the order in which their final commit record appears in the transaction log. XLA filters out records associated with changes to the database that have not yet committed.

If a change is made but then rolled back, XLA does not deliver the records for the rolled back transaction to the application.

Consider the transaction log illustrated in [Figure 3-1](#) and the example that follows, illustrating most of these XLA concepts.

**Figure 3-1 Records Extracted from the Transaction Log**



In this example, the transaction log contains the following records:

- CT<sub>1</sub> - Application C updates row 1 of table W with value 7.7.
- BT<sub>1</sub> - Application B updates row 3 of table X with value 2.
- CT<sub>2</sub> - Application C updates row 9 of table W with value 5.6.
- BT<sub>2</sub> - Application B updates row 2 of table Y with value "XYZ".
- AT<sub>1</sub> - Application A updates row 1 of table Z with value 3.
- AT<sub>2</sub> - Application A updates row 3 of table Z with value 4.
- BT<sub>3</sub> - Application B commits its transaction.
- AT<sub>3</sub> - Application A rolls back its transaction.
- CT<sub>3</sub> - Application C commits its transaction.

An XLA application that is set up to detect changes to tables W, Y, and Z would see the following:

- BT<sub>2</sub> and BT<sub>3</sub> - Update row 2 of table Y with value "XYZ" and commit.
- CT<sub>1</sub> - Update row 1 of table W with value 7.7.
- CT<sub>2</sub> and CT<sub>3</sub> - Update row 9 of table W with value 5.6 and commit.

This example demonstrates the following:

- Transaction records for application B and application C all appear.
- Though the records for application C begin to appear in the transaction log before those for application B, the commit for application B (BT<sub>3</sub>) appears in the transaction log before the commit for application C (CT<sub>3</sub>). As a result, the records for application B are returned to the XLA application ahead of those for application C.
- The application B update to table X (BT<sub>1</sub>) is not presented because XLA is not set up to detect changes to table X.
- The application A updates to table Z (AT<sub>1</sub> and AT<sub>2</sub>) are never presented because it did not commit and was rolled back (AT<sub>3</sub>).

## XLA and Materialized Views

You can use XLA to track changes to both tables and materialized views.

A materialized view provides a single source from which you can track changes to selected rows and columns in multiple detail tables. Without a materialized view, the XLA application

would have to monitor and filter the update records from all of the detail tables, including records reflecting updates to rows and columns of no interest to the application.

In general, there are no operational differences between the XLA mechanisms used to track changes to a table or a materialized view.

For more information about materialized views, see the following:

- [CREATE MATERIALIZED VIEW](#) in *Oracle TimesTen In-Memory Database SQL Reference*
- [Understanding Materialized Views](#) in *Oracle TimesTen In-Memory Database Operations Guide*

## XLA Bookmarks

This section describes XLA bookmarks.

- [How Bookmarks Work](#)
- [Replicated Bookmarks](#)
- [XLA Bookmarks and Transaction Log Holds](#)

## How Bookmarks Work

An XLA bookmark marks the read position of an XLA subscriber application in the transaction log. Bookmarks facilitate durable subscriptions, enabling an application to disconnect from a topic and then reconnect to continue receiving updates where it left off.

When you create a message consumer for XLA, you always use a durable `TopicSubscriber`. The subscription identifier you specify when you create the subscriber is used as the XLA bookmark name. When you use the `ttXlaSubscribe` and `ttXlaUnsubscribe` built-in procedures through JDBC to start and stop the XLA subscription for a table, you explicitly specify the name of the bookmark to be used.

Bookmarks are reset to the last read position whenever an acknowledgment is received. For more information about how update messages are acknowledged, see the [XLA Update Acknowledgments](#).

You can remove a durable subscription by calling `unsubscribe()` on the `JMS Session` object. This deletes the corresponding XLA bookmark and forces a new subscription to be created when you reconnect. For more information see "[Delete Bookmarks](#)".

A bookmark subscription cannot be altered when it is in use. To alter a subscription, you must close the message consumer, alter the subscription using `ttXlaSubscribe` and `ttXlaUnsubscribe`, and open the message consumer.

### Note

You can also use the `ttXlaBookmarkCreate` TimesTen built-in procedure to create bookmarks. See `ttXlaBookmarkCreate` in *Oracle TimesTen In-Memory Database Reference*.

## Replicated Bookmarks

If you are using an active standby pair replication scheme, you have the option of using *replicated bookmarks*, according to the `replicatedBookmark` attribute of the `<topic>` element in the `jmsxla.xml` file.

See [JMS/XLA Configuration File and Topics](#).

For a replicated bookmark, operations on the bookmark are replicated to the standby database as appropriate, assuming there is suitable write privilege for the standby. This enables more efficient recovery of your bookmark positions if a failover occurs.

When you use replicated bookmarks, steps must be taken in the following order:

1. Create the active standby pair replication scheme. (This is accomplished by the `create active standby pair` operation, or by the `ttCWAdmin -create` command in a Clusterware-managed environment.)
2. Create the bookmarks.
3. Subscribe the bookmarks.
4. Start the active standby pair, at which time duplication to the standby occurs and replication begins. (This is accomplished by the `ttRepAdmin -duplicate` command, or by the `ttCWAdmin -start` command in a Clusterware-managed environment.)

### Note

- Alternatively, if you use `ttXlaBookmarkCreate` to create a bookmark, that function has a parameter for specifying a replicated bookmark.
- If you specify replicated bookmarks in the JMS/XLA configuration file, JMS/XLA will create and subscribe to the bookmarks when the application is started. (Also see [JMS/XLA Configuration File and Topics](#).)

Be aware of the following usage notes:

- The position of the bookmark in the standby database is very close to that of the bookmark in the active database; however, because the replication of acknowledge operations is asynchronous, you may see a small window of duplicate updates when there is a failover, depending on how often acknowledge operations are performed.
- It is permissible to drop the active standby pair scheme while replicated bookmarks exist. The bookmarks of course cease to be replicated at that point, but are not deleted. If you subsequently re-enable the active standby pair scheme, these bookmarks are automatically added to the scheme.
- You cannot delete replicated bookmarks while the replication agent is running.
- You can only read and acknowledge a replicated bookmark in the active database. Each time you acknowledge a replicated bookmark, the acknowledge operation is asynchronously replicated to the standby database.

## XLA Bookmarks and Transaction Log Holds

You should be aware that when XLA is in use, there is a hold on TimesTen transaction log files until the XLA bookmark advances.

The hold prevents transaction log files from being purged until XLA can confirm it no longer needs them. If a bookmark becomes stuck, which can occur if an XLA application terminates unexpectedly or disconnects without first deleting its bookmark or disabling change tracking, the log hold persists and there may be an excessive accumulation of transaction log files. This accumulation may result in file system space being filled.

See *Monitoring Accumulation of Transaction Log Files* in *Oracle TimesTen In-Memory Database Operations Guide*.

## JMS/XLA Configuration File and Topics

To connect to XLA, establish a connection to a `JMS Topic` object that corresponds to a particular database. The JMS/XLA configuration file provides the mapping between topic names and databases.

You can specify a replicated bookmark by setting `replicatedBookmark="yes"` in the `<topic>` element when you specify the topic. The default setting is "no". Also see [XLA Bookmarks](#).

By default, JMS/XLA looks for a configuration file named `jmsxla.xml` in the current working directory. If you want to use another name or location for the file, you must specify it as part of the environment variable in the `InitialContext` class and add the location to the classpath.

The following code specifies the configuration file as part of the environment variable in the `InitialContext` class.

```
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
    "com.timesten.dataserver.jmsxla.SimpleInitialContextFactory");
env.put(XlaConstants.CONFIG_FILE_NAME, "/newlocation.xml");
InitialContext ic = new InitialContext(env);
```

The JMS/XLA API uses the class loader to locate the JMS/XLA configuration file if `XlaConstants.CONFIG_FILE_NAME` is set. In this example, the JMS/XLA API searches for the `newlocation.xml` file in the top directory in both the location specified in the `CLASSPATH` environment variable and in the JAR files specified in the `CLASSPATH` variable.

The JMS/XLA configuration file can also be located in subdirectories, as follows:

```
env.put(XlaConstants.CONFIG_FILE_NAME,
    "/com/mycompany/myapplication/deepinside.xml");
```

In this case, the JMS/XLA API searches for the `deepinside.xml` file in the `com/mycompany/myapplication` subdirectory in both the location specified in the `CLASSPATH` environment variable and in the JAR files specified in the `CLASSPATH` variable.

The JMS/XLA API uses the first configuration file that it finds.

A topic definition in the configuration file consists of a name, a connection string, and a prefetch value that specifies how many updates to retrieve at a time. For example, this configuration maps the `DemoDataStore` topic to the `TestDB` DSN:

```
<xlaconfig>
  <topics>
    <topic name="DemoDataStore"
      connectionString="DSN=TestDB"
      xlaPrefetch="100" />
```

```
</topics>
</xlaconfig>
```

A topic definition can also specify whether a replicated bookmark should be used. The following repeats the preceding example, but with a replicated bookmark.

```
<xlaconfig>
  <topics>
    <topic name="DemoDataStore"
      connectionString="DSN=TestDB"
      xlaPrefetch="100" replicatedBookmark="yes" />
  </topics>
</xlaconfig>
```

## XLA Updates

Applications receive XLA updates as JMS `MapMessage` objects. A `MapMessage` object contains a set of typed name and value pairs that correspond to the fields in an XLA update header.

You can access the message fields using the `MapMessage` getter methods. The `getMapNames()` method returns an `Enumeration` object that contains the names of all of the fields in the message. You can retrieve individual fields from the message by name. All reserved field names begin with two underscores, for example `__TYPE`.

All update messages have a `__TYPE` field that indicates what type of update the message contains. The types are specified as integer values. As a convenience, you can use the constants defined in `com.timesten.dataserver.jmsxla.XlaConstants` to compare against the integer types. The supported types are described in [Table 3-1](#).

**Table 3-1 XLA Update Types**

Update type	Description
INSERT	A row has been added.
UPDATE	A row has been modified.
DELETE	A row has been removed.
COMMIT_ONLY	A transaction has been committed.
CREATE_TABLE	A table has been created.
DROP_TABLE	A table has been dropped.
CREATE_INDEX	An index has been created.
DROP_INDEX	An index has been dropped.
ADD_COLUMNS	New columns have been added to the table.
DROP_COLUMNS	Columns have been removed from the table.
CREATE_VIEW	A materialized view has been created.
DROP_VIEW	A materialized view has been dropped.
CREATE_SEQ	A sequence has been created.
DROP_SEQ	A sequence has been dropped.
CREATE_SYNONYM	A synonym has been created.

**Table 3-1 (Cont.) XLA Update Types**

Update type	Description
DROP_SYNONYM	A synonym has been dropped.
TRUNCATE	All rows in the table have been deleted.

For more information about the contents of an XLA update message, see "[JMS/XLA MapMessage Contents](#)".

## XLA Update Acknowledgments

This section discusses acknowledgment of updates in JMS/XLA.

- [XLA Acknowledgment Mechanism](#)
- [XLA Acknowledgment Modes](#)
- [Prefetching Updates](#)
- [Acknowledging Updates](#)

## XLA Acknowledgment Mechanism

The XLA acknowledgment mechanism is designed to ensure that an application has not only received a message, but has successfully processed it.

Acknowledging an update permanently resets the application XLA bookmark to the last record that was read. This prevents previously returned records from being reread, ensuring that an application receives only new batches of records if the bookmark is reused when an application reconnects to XLA.

JMS/XLA can automatically acknowledge XLA update messages, or applications can choose to acknowledge messages explicitly. You specify how updates are to be acknowledged when you create the `Session` object.

## XLA Acknowledgment Modes

JMS/XLA supports three acknowledgment modes.

- **AUTO\_ACKNOWLEDGE:** In this mode, updates are automatically acknowledged as you receive them. Each message is delivered only once. Duplicate messages are not sent, so messages might be lost if there is an application failure. Messages are always delivered and acknowledged individually, so JMS/XLA does not prefetch multiple records. The `xlaprefetch` attribute in the topic is ignored.
- **DUPS\_OK\_ACKNOWLEDGE:** In this mode, updates are automatically acknowledged, but duplicate messages might be delivered when there is an application failure. JMS/XLA prefetches records according to the `xlaprefetch` attribute specified for the topic and sends an acknowledgment when the last record in a prefetched block is read. If the application fails before reading all of the prefetched records, all of the records in the block are presented to the application it restarts.

See [JMS/XLA Configuration File and Topics](#) for examples setting `xlaprefetch`.

- **CLIENT\_ACKNOWLEDGE:** In this mode, applications are responsible for acknowledging receipt of update messages by calling `acknowledge()` on the `MapMessage` instance. JMS/XLA prefetches records according to the `xlaprefetch` attribute specified for the topic.

The following example sets the acknowledgment mode:

```
Session session = connection.createSession (false,  
Session.CLIENT_ACKNOWLEDGE);
```

Also see [Reduce Frequency of Update Acknowledgments](#).

## Prefetching Updates

Prefetching multiple update records at a time is more efficient than obtaining each update record from XLA individually.

Because updates are not prefetched when you use `AUTO_ACKNOWLEDGE` mode, it can be slower than the other modes. If possible, you should design your application to tolerate duplicate updates so you can use `DUPS_OK_ACKNOWLEDGE`, or explicitly acknowledge updates. Explicitly acknowledging updates usually yields the best performance, if you can avoid acknowledging each message individually.

## Acknowledging Updates

To explicitly acknowledge an XLA update, call `acknowledge()` on the update message.

Acknowledging a message implicitly acknowledges all previous messages. Typically, you receive and process multiple update messages between acknowledgments. If you are using the `CLIENT_ACKNOWLEDGE` mode and intend to reuse a durable subscription in the future, you should call `acknowledge()` to reset the bookmark to the last-read position before exiting.

## XLA System Privilege

An XLA user must have the XLA system privilege.

- Any XLA functionality requires the system privilege XLA. This includes connecting to TimesTen (which also requires the `CREATE SESSION` privilege) as an XLA reader and executing the TimesTen XLA built-in procedures `ttXlaBookmarkCreate`, `ttXlaBookmarkDelete`, `ttXlaSubscribe`, and `ttXlaUnsubscribe`, all of which are documented in *Built-In Procedures in Oracle TimesTen In-Memory Database Reference*.
- A user with the XLA privilege has capabilities equivalent to the `SELECT ANY TABLE`, `SELECT ANY VIEW`, and `SELECT ANY SEQUENCE` system privileges.

## XLA Limitations

Be aware of XLA limitations when you use TimesTen JMS/XLA.

- JMS/XLA is available on all platforms supported by TimesTen. However, XLA does not support data transfer between different platforms.
- JMS/XLA support for LOBs is limited. See [Monitoring Tables for Updates](#).
- JMS/XLA does not support applications linked with a driver manager library or the client/server library.
- An XLA reader cannot subscribe to a table that uses in-memory column-based compression.
- For autorefresh cache groups, the change-tracking trigger on Oracle Database does not have column-level resolution. (To have that would be very expensive.) Therefore the

autorefresh feature updates all the columns in the row, and XLA can only report that all the columns have changed, even if data did not actually change in all columns.

## JMS/XLA and Oracle GDK Dependency

The JMS/XLA API uses `orai18n.jar`, part of the Oracle Globalization Development Kit (GDK) for translating from the database character set specified by the `DatabaseCharacterSet` attribute to UTF-16 encoding.

The JMS/XLA API supports a specific version of the GDK with each TimesTen release. If JMS/XLA finds other versions of the GDK loaded in the JVM, it displays a severe warning and continues processing. You can find out the GDK version supported by JMS/XLA by entering the following commands:

```
$ cd timesten_home/install/lib
$ java -cp ./orai18n.jar oracle.i18n.util.GDKOracleMetaData -version
```

Also see [Setting the Module Path or Class Path for Java Development](#).

### Note

The path `timesten_home/install` is a symbolic link to `installation_dir`.

## JMS/XLA Usage and Functionality

There are key aspects when using JMS/XLA.

- [Connecting to XLA](#)
- [Monitoring Tables for Updates](#)
- [Receiving and Processing Updates](#)
- [Terminating a JMS/XLA Application](#)

## Connecting to XLA

To connect to XLA so you can receive updates:

1. Create the appropriate JMS connection factory depending on whether your application uses the JMS API version 1.1 or the Jakarta Messaging API version 1.1 using the property `Context.INITIAL_CONTEXT_FACTORY` when creating the initial context.
  - For Jakarta Messaging API version 1.1, use the value `com.timesten.dataserver.jakartajmsxla.SimpleInitialContextFactory`.
  - For JMS API version 1.1, use the value `com.timesten.dataserver.jmsxla.SimpleInitialContextFactory`.

If the property is not specified, then TimesTen JMS/XLA will default to the JMS API version 1.1 provider unless `download-directory/jakarta.jms-api-3.1.0.jar` or `download-directory/jakarta.jms-api-3.0.0.jar` is present in the `CLASSPATH`, in which case, TimesTen JMS/XLA will use the Jakarta Messaging API version 1.1 provider. The JARs `jakarta.jms-api-3.1.0.jar` and `jakarta.jms-api-3.0.0.jar` contains the class

`jakarta.jms.TopicConnectionFactory` required for the Jakarta Messaging API version 1.1 provider. See [Setting the Module Path or Class Path for Java Development](#).

2. Use the JMS connection factory to create a connection.
3. Use the connection to establish a session.
4. When you are ready to start processing updates, call `start()` on the connection to enable message dispatching.

The following example is from the `asyncJMS2` sample program for Jakarta JMS in the TimesTen Classic Quick Start sample application. See [About TimesTen Quick Start and Sample Applications](#).

```

/** JMS connection */
private jakarta.jms.TopicConnection connection;
...
//jakarta jms

Properties props = new Properties();

props.setProperty(Context.INITIAL_CONTEXT_FACTORY,
    "com.timesten.dataserver.jakartajmsxla.SimpleInitialContextFactory");

// get Connection
Context messaging = new InitialContext(props);
TopicConnectionFactory connectionFactory =
(TopicConnectionFactory)messaging.lookup("TopicConnectionFactory");
connection = connectionFactory.createTopicConnection();
connection.start();

// get Session
log("create session");
session = connection.createTopicSession(false,
Session.AUTO_ACKNOWLEDGE);
...

```

For JavaX JMS, refer to the `asyncJMS` sample program instead in the TimesTen Classic Quick Start.

## Monitoring Tables for Updates

Before you can start receiving updates, you must inform XLA which tables you want to monitor for changes.

To subscribe to changes and turn on XLA publishing for a table, call the `ttXlaSubscribe` built-in procedure through JDBC.

When you use `ttXlaSubscribe` to enable XLA publishing for a table, you must specify parameters for the name of the table and the name of the bookmark that are used to track the table:

```
ttXlaSubscribe(user.table, mybookmark)
```

For example, call `ttXlaSubscribe` by the JDBC `CallableStatement` interface:

```
Connection con;

CallableStatement cStmt;
...
cStmt = con.prepareCall("{call ttXlaSubscribe(user.table, mybookmark)}");
cStmt.execute();
```

Use `ttXlaUnsubscribe` to unsubscribe from the table during shutdown. For more information, see [Unsubscribe from a Table](#).

The application can verify table subscriptions by checking the `SYS.XLASUBSCRIPTIONS` system table.

For more information about using TimesTen built-in procedures in a Java application, see [Using CALL to Execute Procedures and Functions](#).

### Note

LOB support in JMS/XLA is limited, as follows:

- You can subscribe to tables containing LOB columns, but information about the LOB value itself is unavailable.
- Columns containing LOBs are reported as empty (zero length) or null (if the value is actually `NULL`). In this way, you can tell the difference between a null column and a non-null column.

See the next section, [Receiving and Processing Updates](#), for additional notes.

## Receiving and Processing Updates

You can receive XLA updates either synchronously or asynchronously.

To receive and process updates for a topic synchronously, perform the following tasks.

1. Create a durable `TopicSubscriber` instance to subscribe to a topic.
2. Call `receive()` or `receiveNoWait()` on your subscriber to get the next available update.
3. Process the returned `MapMessage` instance.

To receive and process updates for a topic asynchronously, perform the following tasks.

1. Create a `MessageListener` instance to process the updates.
2. Create a durable `TopicSubscriber` instance to subscribe to a topic.
3. Register the `MessageListener` with the `TopicSubscriber`.
4. Start the connection.

**Note**

You may miss messages if you do not register the `MessageListener` before you start the connection. If the connection is already started, stop the connection, register the `MessageListener`, then start the connection.

5. Wait for messages to arrive. You can call the `Object` method `wait()` to wait for messages if your application does not have to do anything else in its main thread.

When an update is published, the `MessageListener` method `onMessage()` is called and the message is passed in as a `MapMessage` instance.

The application can verify table subscriptions by checking the `SYS.XLASUBSCRIPTIONS` system table.

**Note**

LOB support in XLA is limited. You can access LOB fields in update messages using the `MapMessage` method `getBytes()` for BLOB fields or `getString()` for CLOB or NCLOB fields; however, these fields contain zero-length data (or null data if the value is actually `NULL`).

The following example, from the `asyncJMS` TimesTen Classic Quick Start sample application, uses a listener to process updates asynchronously.

```
MyListener myListener = new MyListener(outStream);

outStream.println("Creating consumer for topic " + topic);
Topic xlaTopic = session.createTopic(topic);
bookmark = "bookmark";
TopicSubscriber subscriber = session.createDurableSubscriber(xlaTopic,
bookmark);

// After setMessageListener() has been called, myListener's onMessage
// method is called for each message received.
subscriber.setMessageListener(myListener);
```

Note that `bookmark` must already exist. You can use JDBC and the `ttXlaBookmarkCreate` built-in procedure to create a bookmark. Also, the `TopicSubscriber` must be a durable subscriber. XLA connections are designed to be durable. XLA bookmarks make it possible to disconnect from a topic and then reconnect to start receiving updates where you left off. The string you pass in as the subscription identifier when you create a durable subscriber is used as the XLA bookmark name.

You can call `unsubscribe()` on the JMS `TopicSession` to delete the XLA bookmark used by the subscriber when the application shuts down. This causes a new bookmark to be created when the application is restarted.

When you receive an update, you can use the `MapMessage` getter methods to extract information from the message and then perform whatever processing your application requires. The TimesTen `XlaConstants` class defines constants for the update types and special message fields for use in processing XLA update messages.

The first step is typically to determine what type of update the message contains. You can use the `MapMessage` method `getInt()` to get the contents of the `__TYPE` field, and compare the value against the numeric constants defined in the `XlaConstants` class.

In the next example, from the `asyncJMS TimesTen Classic Quick Start` sample application, the method `onMessage()` extracts the update type from the `MapMessage` object and displays the action that the update signifies.

```
public void onMessage(Message message)
{
    MapMessage mapMessage = (MapMessage)message;
    String messageType = null;
    /* Standard output stream */
    private static PrintStream outputStream = System.out;

    if (message == null)
    {
        errStream.println("MyListener: update message is null");
        return ;
    }

    try
    {
        outputStream.println();
        outputStream.println("onMessage: got a " + mapMessage.getJMSType() + "
message");

        // Get the type of event (insert, update, delete, drop table, etc.).
        int type = mapMessage.getInt(XlaConstants.TYPE_FIELD);
        if (type == XlaConstants.INSERT)
        {
            outputStream.println("A row was inserted.");
        }
        else if (type == XlaConstants.UPDATE)
        {
            outputStream.println("A row was updated.");
        }
        else if (type == XlaConstants.DELETE)
        {
            outputStream.println("A row was deleted.");
        }
        else
        {
            // Messages are also received for DDL events such as CREATE TABLE.
            // This program processes INSERT, UPDATE, and DELETE events,
            // and ignores the DDL events.
            return ;
        }
    }
    ...
}
...
}
```

When you know what type of message you have received, you can process the message according to the application's needs. To get a list of all of the fields in a message, you can call

the `MapMessage` method `getMapNames()`. You can retrieve individual fields from the message by name.

The following example, from the `asyncJMS TimesTen Classic Quick Start` sample application, extracts the column values from insert, update, and delete messages using the column names.

```
/* Standard output stream */
private static PrintStream outputStream = System.out;
...
if (type == XlaConstants.INSERT
    || type == XlaConstants.UPDATE
    || type == XlaConstants.DELETE)
{

    // Get the column values from the message.
    int cust_num = mapMessage.getInt("cust_num");
    String region = mapMessage.getString("region");
    String name = mapMessage.getString("name");
    String address = mapMessage.getString("address");

    outputStream.println("New Column Values:");
    outputStream.println("cust_num=" + cust_num);
    outputStream.println("region=" + region);
    outputStream.println("name=" + name);
    outputStream.println("address=" + address);
}
```

For detailed information about the contents of XLA update messages, see [JMS/XLA MapMessage Contents](#). For information about how TimesTen column types map to JMS data types and the getter methods used to retrieve the column values, see [Data Type Support](#).

## Terminating a JMS/XLA Application

When the application has finished reading from the transaction log, there are steps to enable the application to gracefully exit.

- [Close the Connection](#)
- [Delete Bookmarks](#)
- [Unsubscribe from a Table](#)

### Close the Connection

To close the connection to XLA, call `close()` on the `Connection` object.

After a connection has been closed, any attempt to use it, its sessions, or its subscribers results in an `IllegalStateException` error. You can continue to use messages received through the connection, but you cannot call the `acknowledge()` method on the received message after the connection is closed.

### Delete Bookmarks

Deleting XLA bookmarks during shutdown is optional. Deleting a bookmark enables the file system space associated with any unread update records in the transaction log to be freed.

If you do not delete the bookmark, it can be reused by a durable subscriber. If the bookmark is available when a durable subscriber reconnects, the subscriber receives all unacknowledged updates published since the previous connection was terminated. Keep in mind that when a bookmark exists with no application reading from it, the transaction log continues to grow and the amount of file system space consumed by your database increases.

To delete a bookmark, you can simply call `unsubscribe()` on the JMS Session, which invokes the `ttXlaBookmarkDelete` built-in procedure to remove the XLA bookmark.

**Note**

You cannot delete replicated bookmarks while the replication agent is running.

## Unsubscribe from a Table

To turn off XLA publishing for a table, use the `ttXlaUnsubscribe` built-in procedure. If you use `ttXlaSubscribe` to enable XLA publishing for a table, use `ttXlaUnsubscribe` to unsubscribe from the table when shutting down your application.

**Note**

If you want to drop a table, you must unsubscribe from it first.

When you unsubscribe from a table, specify the name of the table and the name of the bookmark used to track the table:

```
ttXlaUnsubscribe(user.table, mybookmark)
```

The following example calls `ttXlaUnSubscribe` through a `CallableStatement` object.

```
Connection con;  
CallableStatement cStmt;  
...  
cStmt = con.prepareCall("{call ttXlaUnSubscribe(user.table, mybookmark)}");  
cStmt.execute();
```

See [Using CALL to Execute Procedures and Functions](#).

## JMS/XLA as a Replication Mechanism

In the event that TimesTen replication is not sufficient for your needs, you can use JMS/XLA instead.

- [About Using JMS/XLA as a Replication Mechanism](#)
- [Applying JMS/XLA Messages to a Target Database](#)
- [TargetDataStore Error Recovery](#)

## About Using JMS/XLA as a Replication Mechanism

TimesTen replication is sufficient for most customer needs; however, it is also possible to use JMS/XLA to replicate updates from one database to another.

Implementing your own replication scheme on top of JMS/XLA in this way is fairly complicated, but can be considered if TimesTen replication is not feasible for some reason.

Overview of TimesTen Replication in the *Oracle TimesTen In-Memory Database Replication Guide* describes TimesTen replication.

## Applying JMS/XLA Messages to a Target Database

The source database generates JMS/XLA messages. To apply the messages to a target database, you must extract the XLA descriptor from them.

Use the `MapMessage` interface to extract the update descriptor:

```
MapMessage message;
/*
 *...other code
 */
try {
    byte[] updateMessage=
        mapMessage.getBytes(XlaConstants.UPDATE_DESCRIPTOR_FIELD);
}
catch (JMSEException jex){
/*
 *...other code
 */
}
```

The target database may reside on a different system from the source database. The update descriptor is returned as a byte array and can be serialized for network transmission.

You must create a target database object that represents the target database so you can apply the objects from the source database. You can create a target database object named `myTargetDataStore` as an instance of the `TargetDataStoreImpl` class. For example:

```
TargetDataStore myTargetDataStore=
    new TargetDataStoreImpl("DSN=sampled");
```

Apply messages to `myTargetDataStore` by using the `TargetDataStore` method `apply()`. For example:

```
myTargetDataStore.apply(updateDescriptor);
```

By default, TimesTen checks for conflicts on the target database before applying the update. If the target database has information that is later than the update, `TargetDataStore` throws an exception. If you do not want TimesTen to check for conflicts, use the `TargetDataStore` method `setUpdateConflictCheckFlag()` to change the behavior.

By default, TimesTen commits the update to the database based on commit flags and transaction boundaries contained in the update descriptor. If you want the application to

perform manual commits instead, use the `setAutoCommitFlag()` method to change the autocommit flag. To perform a manual commit on `myTargetDataStore`, use the following command:

```
myTargetDataStore.commit();
```

You can perform a rollback if errors occur during the application of the update. Use the following command for `myTargetDataStore`:

```
myTargetDataStore.rollback();
```

Close `myTargetDataStore` by using the following command:

```
myTargetDataStore.close();
```

See [JMS/XLA Replication API](#).

## TargetDataStore Error Recovery

Invoking `TargetDataStore` can yield transient and permanent errors.

This section discusses error recovery.

`TargetDataStore` methods return a nonzero value when transient errors occur. The application can retry the operation and is responsible for monitoring update descriptors that must be reapplied. For more information about transient XLA errors, see *XLA Error Handling in Oracle TimesTen In-Memory Database C Developer's Guide*.

`TargetDataStore` methods return a `JMSException` object for permanent errors. If the application receives a permanent error, it should verify that the database is valid. If the database is invalid, the target database object should be closed and a new one should be created. Other types of permanent errors may require manual intervention.

The following example shows how to recover errors from a `TargetDataStore` object.

```
TargetDataStore theTargetDataStore;  
byte[] updateDescriptor;  
int rc;  
  
// Other code  
try {  
    ...  
    if ( (rc = theTargetDataStore.apply(updateDescriptor) ) == 0 ) {  
        // Apply successful.  
    }  
    else {  
        // Transient error. Retry later.  
    }  
}  
catch (JMSException jex) {  
    if (theTargetDataStore.isDataStoreValid() ) {  
        // Database valid; permanent error that may need Administrator  
        intervention.  
    }  
    else {
```

```
    try {
        theTargetDataStore.close();
    }
    catch (JMSEException closeEx) {
        // Close errors are not usual. This may need Administrator intervention.
    }
}
```

# 4

## Distributed Transaction Processing: JTA

This chapter describes the implementation of the Java Transaction API (JTA) for TimesTen Classic.

- [Overview of JTA](#)
- [JTA Functionality in TimesTen](#)
- [Using the JTA API](#)

### Overview of JTA

This is a brief overview of XA concepts.

- [About TimesTen JTA](#)
- [X/Open DTP Model](#)
- [Two-Phase Commit](#)

### About TimesTen JTA

The TimesTen implementation of the JTA interfaces, supported only in TimesTen Classic, is intended to enable Java applications, application servers, and transaction managers to use TimesTen resource managers in distributed transaction processing (DTP) environments.

The TimesTen implementation is supported for use with the Oracle WebLogic Server.

Refer to the following locations for additional information.

- General information about JTA:  
<https://www.oracle.com/java/technologies/jta.html>
- Oracle WebLogic information and documentation:  
<https://www.oracle.com/middleware/technologies/weblogic.html>

As TimesTen JTA is built on top of the TimesTen implementation of the X/Open XA standard, much of the discussion in this chapter is in terms of underlying XA features. You can also refer to Distributed Transaction Processing: XA in *Oracle TimesTen In-Memory Database C Developer's Guide*.

#### ✓ Tip

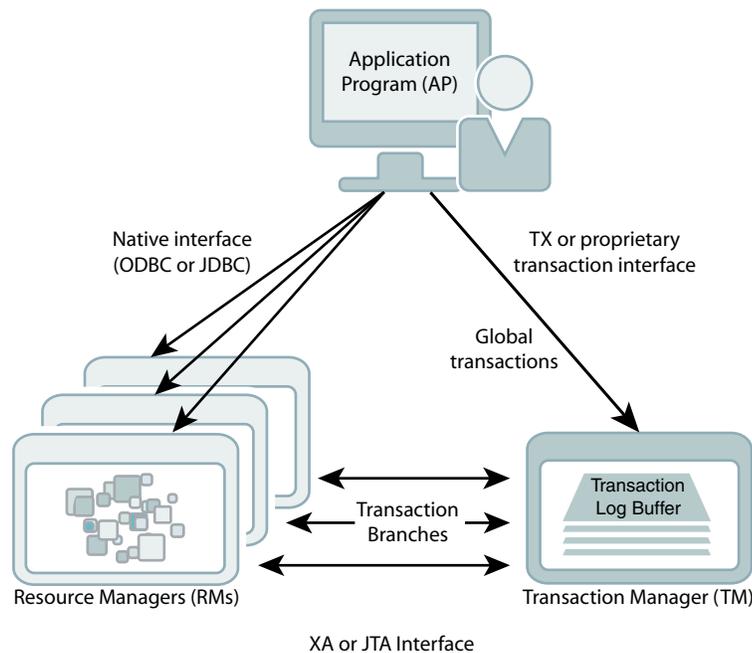
- The TimesTen XA implementation does not work with cache. The start of any XA transaction fails if the cache agent is running.
- You cannot execute an XA transaction if replication is enabled.
- Do not execute DDL statements within an XA transaction.

## X/Open DTP Model

There are interfaces defined by the X/Open DTP model.

[Figure 4-1](#) illustrates the interfaces defined by the X/Open DTP model.

**Figure 4-1 Distributed Transaction Processing Model**



The TX interface is what applications use to communicate with a transaction manager. The figure shows an application communicating global transactions to the transaction manager. In the DTP model, the transaction manager breaks each global transaction down into multiple branches and distributes them to separate resource managers for service. It uses the JTA interface to coordinate each transaction branch with the appropriate resource manager.

In the context of TimesTen JTA, the resource managers can be a collection of TimesTen databases, or databases in combination with other commercial databases that support JTA.

Global transaction control provided by the TX and JTA interfaces is distinct from local transaction control provided by the native JDBC interface. It is generally best to maintain separate connections for local and global transactions. Applications can obtain a connection handle to a TimesTen resource manager to initiate both local and global transactions over the same connection.

## Two-Phase Commit

In a JTA implementation, the transaction manager commits the distributed branches of a global transaction by using a two-phase commit protocol.

1. In phase 1, the transaction manager directs each resource manager to prepare to commit, which is to verify and guarantee it can commit its respective branch of the global transaction. If a resource manager cannot commit its branch, the transaction manager rolls back the entire transaction in phase 2.

2. In phase 2, the transaction manager either directs each resource manager to commit its branch or, if a resource manager reported it was unable to commit in phase 1, rolls back the global transaction.

Note the following optimizations.

- If a global transaction is determined by the transaction manager to have involved only one branch, it skips phase 1 and commits the transaction in phase 2.
- If a global transaction branch is read-only, where it does not generate any transaction log records, the transaction manager commits the branch in phase 1 and skips phase 2 for that branch.

**Note**

The transaction manager considers the global transaction committed if and only if all branches successfully commit.

## JTA Functionality in TimesTen

There are things you should consider when using JTA in TimesTen.

- [TimesTen Database Requirements for XA](#)
- [Global Transaction Recovery in TimesTen](#)
- [XA Error Handling in TimesTen](#)

## TimesTen Database Requirements for XA

To guarantee global transaction consistency, TimesTen XA transaction branches must be durable.

The TimesTen implementation of the `xa_prepare()`, `xa_rollback()`, and `xa_commit()` functions log their actions to the file system, regardless of the value set in the `DurableCommits` general connection attribute or by the `ttDurableCommit` built-in procedure. If you must recover from a failure, both the resource manager and the TimesTen transaction manager have a consistent view of which transaction branches were active in a prepared state at the time of failure.

## Global Transaction Recovery in TimesTen

When a database is loaded from the file system to recover after a failure or unexpected termination, any global transactions that were prepared but not committed are left pending, or in doubt. Processing is not enabled until the disposition of all in-doubt transactions has been resolved.

After connection and recovery are complete, TimesTen checks for in-doubt transactions. If there are no in-doubt transactions, operations can proceed. If there are in-doubt transactions, other connections may be created, but virtually all operations are prohibited on those connections until the in-doubt transactions are resolved. Any other JDBC calls result in the following error:

```
Error 11035 - "In-doubt transactions awaiting resolution in recovery must be resolved first"
```

The list of in-doubt transactions can be retrieved through the XA implementation of `xa_recover()`, then dealt with through the XA call `xa_commit()`, `xa_rollback()`, or `xa_forget()`, as appropriate. After all the in-doubt transactions are cleared, operations can proceed.

This scheme should be adequate for systems that operate strictly under control of the transaction manager, since the first thing the transaction manager should do after connect is to call `xa_recover()`.

If the transaction manager is unavailable or cannot resolve an in-doubt transaction, you can use the `ttXactAdmin` utility `-HCommit` or `-HAbort` option to independently commit or roll back the individual transaction branches. Be aware, however, that these `ttXactAdmin` options require `ADMIN` privilege. See `ttXactAdmin` in *Oracle TimesTen In-Memory Database Reference*.

## XA Error Handling in TimesTen

The XA specification has a limited, strictly defined set of errors that can be returned from XA interface calls. The ODBC `SQLERROR` mechanism returns XA defined errors, along with any additional information.

The TimesTen XA related errors begin at number 11000. Errors 11002 through 11020 correspond to the errors defined by the XA standard. See the Error Help portal.

## Using the JTA API

The TimesTen implementation of JTA provides an API consistent with that specified in the JTA specification.

This section covers the following topics for using the JTA API:

- [Registering a TimesTen DSN with WebLogic](#)
- [Importing Required Packages](#)
- [Creating a TimesTen XAConnection Object](#)
- [Creating XAResource and Connection Objects](#)

## Registering a TimesTen DSN with WebLogic

Regarding how to register a TimesTen DSN with WebLogic, information on configuring TimesTen for application servers and object-relational mapping frameworks is available in the TimesTen Classic Quick Start.

See [About TimesTen Quick Start and Sample Applications](#).

## Importing Required Packages

The TimesTen JDBC and XA implementations are available in packages.

```
import com.timesten.jdbc.*;  
import com.timesten.jdbc.xa.*;
```

Your application should also import these standard packages:

```
import java.sql.*;
import javax.sql.*;
import javax.transaction.xa.*;
```

## Creating a TimesTen XAConnection Object

Connections to XA data sources are established through `XADataSource` objects. You can create an `XAConnection` object for your database by using the `TimesTenXADataSource` instance as a connection factory. `TimesTenXADataSource` implements the `javax.sql.XADataSource` interface.

After creating a new `TimesTenXADataSource` instance, use the `setUrl()` method to specify a database.

The URL should look similar to the following.

- For a direct connection: `jdbc:timesten:direct:DSNname`
- For a client connection: `jdbc:timesten:client:DSNname`

You can also optionally use the `setUser()` and `setPassword()` methods to set the ID and password for a specific user.

In the following example, the `TimesTenXADataSource` object is used as a factory to create a TimesTen XA data source object. Then the URL that identifies the TimesTen DSN (`dsn1`), the user name (`myName`), and the password (`password`) are set for this `TimesTenXADataSource` instance. Then the `getXAConnection()` method is used to return a connection to the object, `xaConn`.

```
TimesTenXADataSource xads = new TimesTenXADataSource();

xads.setUrl("jdbc:timesten:direct:dsn1");
xads.setUser("myName");
xads.setPassword("password");

XAConnection xaConn = null;
try {
    xaConn = xads.getXAConnection();
}
catch (SQLException e){
    e.printStackTrace();
    return;
}
```

You can create multiple connections to an XA data source object. This example creates a second connection, `xaConn2`:

```
XAConnection xaConn = null;
XAConnection xaConn2 = null;

try {
    xaConn = xads.getXAConnection();
    xaConn2 = xads.getXAConnection();
}
```

This next example creates two instances of `TimesTenXADataSource` for the databases named `dsn1` and `dsn2`. It then creates a connection for `dsn1` and two connections for `dsn2`.

```
TimesTenXADataSource xads = new TimesTenXADataSource();

xads.setUrl("jdbc:timesten:direct:dsn1");
xads.setUser("myName");
xads.setPassword("password");

XAConnection xaConn1 = null;
XAConnection xaConn2 = null;
XAConnection xaConn3 = null;

try {
    xaConn1 = xads.getXAConnection(); // connect to dsn1
}
catch (SQLException e){
    e.printStackTrace();
    return;
}

xads.setUrl("jdbc:timesten:direct:dsn2");
xads.setUser("myName");
xads.setPassword("password");

try {
    xaConn2 = xads.getXAConnection(); // connect to dsn2
    xaConn3 = xads.getXAConnection(); // connect to dsn2
}
catch (SQLException e){
    e.printStackTrace();
    return;
}
```

 **Note**

Once an `XAConnection` is established, autocommit is turned off.

## Creating XAResource and Connection Objects

After using `getXAConnection()` to obtain an `XAConnection` object, you can use the `XAConnection` method `getXAResource()` to obtain an `XAResource` object, then the `XAConnection` method `getConnection()` to obtain a `Connection` object for the underlying connection.

```
//get an XAResource
XAResource xaRes = null;
try {
    xaRes = xaConn.getXAResource();
}
catch (SQLException e){
    e.printStackTrace();
}
```

```
        return;
    }

    //get an underlying physical Connection
    Connection conn = null;
    try {
        conn = xaConn.getConnection();
    }
    catch (SQLException e){
        e.printStackTrace();
        return;
    }
}
```

From this point, you can use the same connection, `conn`, for both local and global transactions. Be aware of the following, however.

- You must commit or roll back an active local transaction before starting a global transaction. Otherwise you get the `XAException` exception `XAER_OUTSIDE`.
- You must end an active global transaction before initiating a local transaction, otherwise you get a `SQLException`, "Illegal combination of local transaction and global (XA) transaction."

# 5

## Java Application Tuning

This chapter provides tips on how to tune a Java application to run optimally on a TimesTen database.

See TimesTen Database Performance Tuning in *Oracle TimesTen In-Memory Database Operations Guide* for more general tuning tips.

The chapter is organized as follows:

- [Tuning JDBC Applications](#)
- [Tuning JMS/XLA Applications](#)

### Tuning JDBC Applications

There are general principles to consider when tuning JDBC applications for TimesTen.

It includes the following topics:

- [Use Prepared Statement Pooling](#)
- [Use Arrays of Parameters for Batch Execution](#)
- [Bulk Fetch Rows of TimesTen Data](#)
- [Use the ResultSet Method getString\(\) Sparingly](#)
- [Avoid Data Type Conversions](#)
- [Close Connections, Statements and Result Sets](#)
- [Optimize Queries](#)

#### Note

Also see [Working with TimesTen Result Sets: Hints and Restrictions](#) and the notes in [Parameter Binding and Statement Execution](#).

### Use Prepared Statement Pooling

TimesTen supports prepared statement pooling for pooled connections through the TimesTen `ObservableConnectionDS` class.

This is the TimesTen implementation of `ConnectionPoolDataSource`. Note that statement pooling is transparent to an application. Use of the `PreparedStatement` object, including preparing and closing the statement, is no different.

Enable prepared statement pooling and specify the maximum number of statements in the pool by calling the `ObservableConnectionDS` method `setMaxStatements()`. A value of 0, the default, disables prepared statement pooling. Any integer value greater than 0 enables prepared statement pooling with the value taken as the maximum number of statements. Once set, this value should not be changed.

Prepared statements or callable statements are pooled at the time of creation if the pool has not reached its capacity. You can remove a prepared statement or callable statement from the pool by calling `setPoolable(false)` on the statement object. After the statement is closed, it is removed from the pool.

✓ **Tip**

With prepared statement pooling, JDBC considers two statements to be identical if their SQL (including comments) is identical, regardless of other considerations such as optimizer settings. Do not use prepared statement pooling in a scenario where different optimizer hints may be applied to statements that are otherwise identical. In this scenario, a statement execution may result in the use of an identical statement from the pool with an unpredictable optimizer setting.

## Use Arrays of Parameters for Batch Execution

You can improve performance by using groups, referred to as *batches*, of statement executions, calling the `addBatch()` and `executeBatch()` methods for `Statement` or `PreparedStatement` objects.

A batch can consist of a set of `INSERT`, `UPDATE`, `DELETE`, or `MERGE` statements. Statements that return result sets, such as `SELECT` statements, are not allowed in a batch. A SQL statement is added to a batch by calling `addBatch()` on the statement object. The set of SQL statements associated with a batch are executed through the `executeBatch()` method.

For `PreparedStatement` objects, the batch consists of repeated executions of a statement using different input parameter values. For each set of input values, create the batch by using appropriate `setXXX()` calls followed by the `addBatch()` call. The batch is executed by the `executeBatch()` method.

The ideal batch size for any of these database operations varies according to details of the user environment and requires testing and experimentation to determine.

✓ **Tip**

In TimesTen Classic, it is important to use a batch size that is an exact multiple of 256 for inserts, in order to optimize the insert mechanism.

Here is an example of batching statements:

```
// turn off autocommit
conn.setAutoCommit(false);

Statement stmt = conn.createStatement();
stmt.addBatch("INSERT INTO employees VALUES (1000, 'Joe Jones')");
stmt.addBatch("INSERT INTO departments VALUES (260, 'Shoe')");
stmt.addBatch("INSERT INTO emp_dept VALUES (1000, 260)");

// submit a batch of update commands for execution
int[] updateCounts = stmt.executeBatch();
conn.commit ();
```

The following example batches prepared statements:

```
// turn off autocommit
conn.setAutoCommit(false);
// prepare the statement
PreparedStatement stmt = conn.prepareStatement
    ("INSERT INTO employees VALUES (?, ?)");

// first set of parameters
stmt.setInt(1, 2000);
stmt.setString(2, "Kelly Kaufmann");
stmt.addBatch();

// second set of parameters
stmt.setInt(1, 3000);
stmt.setString(2, "Bill Barnes");
stmt.addBatch();

// submit the batch for execution. Check update counts
int[] updateCounts = stmt.executeBatch();
conn.commit ();
```

For either a `Statement` or `PreparedStatement` object, the `executeBatch()` method returns an array of update counts (`updateCounts[]` in the examples above), with one element in the array for each statement execution. The value of each element can be any of the following:

- A number indicating how many rows in the database were affected by the corresponding statement execution
- `SUCCESS_NO_INFO`, indicating the corresponding statement execution was successful, but the number of affected rows is unknown
- `EXECUTE_FAILED`, indicating the corresponding statement execution failed

Once there is a statement execution with `EXECUTE_FAILED` status, no further statement executions are attempted.

For more information about using the JDBC batch update facility, refer to Javadoc for the `java.sql.Statement` interface, particularly information about the `executeBatch()` method, at the following location:

<https://docs.oracle.com/javase/8/docs/api/java/sql/package-summary.html>

#### Note

Associative array parameters are not supported with JDBC batch execution. (See [Working with Associative Arrays](#).)

## Bulk Fetch Rows of TimesTen Data

TimesTen provides an extension that enables an application to fetch multiple rows of data.

For applications that retrieve large amounts of data, fetching multiple rows can increase performance greatly. However, when using Read Committed isolation level, locks are held on

all rows being retrieved until the application has received all the data, decreasing concurrency. See [Fetching Multiple Rows of Data](#).

## Use the ResultSet Method getString() Sparingly

Because Java strings are immutable, the `ResultSet` method `getString()` must allocate space for a new string in addition to translating the underlying C string to a Unicode string, making it a costly call.

In addition, you should not call `getString()` on primitive numeric types, like `byte` or `int`, unless it is absolutely necessary. It is much faster to call `getInt()` on an integer column, for example.

## Avoid Data Type Conversions

TimesTen instruction paths are so short that even small delays due to data conversion can cause a relatively large percentage increase in transaction time.

Use the appropriate `getXXX()` method on a `ResultSet` object for the data type of the data in the underlying database. For example, if the data type of the data is `DOUBLE`, to avoid data conversion in the JDBC driver you should call `getDouble()`. Similarly, use the appropriate `setXXX()` method on the `PreparedStatement` object for the input parameter in an SQL statement. For example, if you are inserting data into a `CHAR` column using a `PreparedStatement`, you should use `setString()`.

## Close Connections, Statements and Result Sets

For better performance, always close JDBC objects such as connection, statement, and result set instances when finished using them.

For example:

```
Connection conn = null;
Statement stmt = null;
ResultSet rs = null;
try {
    // create connections, execute statements, etc.
    // Handle any errors
} catch (SQLException ex) {
    // See Error Handling.
}
finally {
    // Close JDBC objects such as connections, statements, result sets, etc.
    if (rs != null) {
        try {
            rs.close();
        }
        catch(SQLException finalex) {
            // See Error Handling.
        }
    }
    if (stmt != null) {
        try {
            stmt.close();
        }
    }
}
```

```
        catch(SQLException finalex) {
            // See Error Handling.
        }
    }
    // Always, close the connection to TimesTen
    if (conn != null) {
        try {
            conn.close();
        }
        catch(SQLException finalex) {
            // See Error Handling.
        }
    }
}
```

## Optimize Queries

TimesTen provides the `TimesTenConnection` method `setTtPrefetchClose()` to optimize query performance with a `true` setting.

Refer to [Optimizing Query Performance](#).

## Tuning JMS/XLA Applications

There are specific performance tuning tips for applications that use the JMS/XLA API.

- [Considerations in Tuning JMS/XLA Applications](#)
- [Configure xlaPrefetch Parameter](#)
- [Reduce Frequency of Update Acknowledgments](#)
- [Handling High Event Rates](#)

## Considerations in Tuning JMS/XLA Applications

JMS/XLA has some overhead that makes it slower than using the C XLA API.

In the C API, records are returned to the user in a batch. In the JMS model an object is instantiated and each record is presented one at a time in a callback to the `MessageListener` method `onMessage()`. High performance applications can use some tuning to overcome some of this overhead.

## Configure xlaPrefetch Parameter

The code underlying the JMS layer that reads the transaction log is more efficient if it can fetch as many rows as possible before presenting the object/rows to the user. The amount of prefetching is controlled in the `jmsxla.xml` configuration file with the `xlaPrefetch` parameter.

Set the prefetch count to a large value like 100 or 1000.

## Reduce Frequency of Update Acknowledgments

In JMS/XLA, acknowledging updates moves the bookmark and results in updates to system tables. You can typically improve application performance by waiting until several updates have been detected before issuing the acknowledgment. You can control the acknowledgment frequency in either of the following modes.

See [XLA Update Acknowledgments](#).

- `DUPS_OK_ACKNOWLEDGE`, where JMS/XLA prefetches records according to the `xlaprefetch` setting, and an acknowledgment is automatically sent when the last record in the prefetched block is read.
- `CLIENT_ACKNOWLEDGE`, where you manually call the `acknowledge()` method on the `MapMessage` instance as desired.

The appropriate choice for acknowledgment frequency depends on your application logic. Acknowledging after every 100 updates, for example, has been used successfully. Be aware, however, that there is a trade-off. Acknowledgments affect XLA log holds, and acknowledging too infrequently may result in undesirable log file accumulation. (Also see [XLA Bookmarks and Transaction Log Holds](#).)

 **Note**

In `DUPS_OK_ACKNOWLEDGE` or `CLIENT_ACKNOWLEDGE` mode, the reader application must have some tolerance for seeing the same set of records more than once.

## Handling High Event Rates

The synchronous interface is suitable only for applications with low event rates and for which `AUTO_ACKNOWLEDGE` or `DUPS_OK_ACKNOWLEDGE` acknowledgment modes are acceptable.

Applications that require `CLIENT_ACKNOWLEDGE` acknowledgment mode and applications with high event rates should use the asynchronous interface for receiving updates. They should acknowledge the messages on the callback thread itself if they are using `CLIENT_ACKNOWLEDGE` as acknowledgment mode. See [Receiving and Processing Updates](#).

# 6

## JMS/XLA Reference

Reference information is provided for the JMS/XLA API, which is supported by TimesTen Classic.

- [JMS/XLA MapMessage Contents](#)
- [DML Event Data Formats](#)
- [DDL Event Data Formats](#)
- [Data Type Support](#)
- [JMS Interfaces for Event Handling](#)
- [JMS/XLA Replication API](#)
- [JMS Message Header Fields](#)

### JMS/XLA MapMessage Contents

A `jakarta.jms.MapMessage` instance for Jakarta JMS or `javax.jms.MapMessage` instance for JavaX JMS contains a set of typed name/value pairs corresponding to the fields in an XLA update header, which is published as the C structure `ttXlaUpdateDesc_t`.

The fields contained in a `MapMessage` instance depend on what type of update it is.

### XLA Update Types

Each `MapMessage` instance returned by the JMS/XLA API contains at least one name/value pair, `__TYPE` (with 2 underscores), that identifies the type of update described in the message. The types are specified as integer values.

[Table 6-1](#) lists the supported types.

#### Note

As a convenience, you can use the constants defined in `com.timesten.dataserver.jakartajmsxla.XLAConstants` for Jakarta JMS or `com.timesten.dataserver.jmsxla.XlaConstants` for JavaX JMS to compare against the integer types.

**Table 6-1 XLA Update Types**

Type	Description
<code>ADD_COLUMNS</code>	Indicates that columns have been added.
<code>COMMIT_FIELD</code>	This is the name of the field in a message that contains a commit.
<code>COMMIT_ONLY</code>	Indicates that a commit has occurred.

**Table 6-1 (Cont.) XLA Update Types**

Type	Description
CONTEXT_FIELD	This is the name of the field in a message that contains the context value passed to the <code>ttApplicationContext</code> procedure as a byte array.
CREATE_INDEX	Indicates that an index has been created.
CREATE_SEQ	Indicates that a sequence has been created.
CREATE_SYNONYM	Indicates that a synonym has been created.
CREATE_TABLE	Indicates that a table has been created.
CREATE_VIEW	Indicates that a view has been created.
DELETE	Indicates that a row has been deleted.
DROP_COLUMNS	Indicates that columns have been dropped.
DROP_INDEX	Indicates that an index has been dropped.
DROP_SEQ	Indicates that a sequence has been dropped.
DROP_SYNONYM	Indicates that a synonym has been dropped.
DROP_TABLE	Indicates that a table has been dropped.
DROP_VIEW	Indicates that a view has been dropped.
FIRST_FIELD	This is the name of the field that contains the flag that indicates the first record in a transaction.
INSERT	Indicates that a row has been inserted.
MTYP_FIELD	This is the name of the field in a message that contains type information.
MVER_FIELD	This is the name of the field in a message that contains the transaction log file number of the XLA record.
NULLS_FIELD	This is the name of the field in a message that contains the list of fields that have null values.
REPL_FIELD	This is the name of the field in a message that contains the flag that indicates that the update was applied by replication.
TBLNAME_FIELD	This is the name of the field in a message that contains the table name.
TBLOWNER_FIELD	This is the name of the field in a message that specifies the table owner.
TRUNCATE	Indicates that a table has been truncated.
TYPE_FIELD	This is the name of the field in a message that specifies the message type.
UPDATE	Indicates that a row has been updated.
UPDATE_DESCRIPTOR_FIELD	This is the name of the field that returns a <code>ttXlaUpdateDesc_t</code> structure as a byte array.
UPDATED_COLUMNS_FIELD	This is the name of the field in a message that contains the list of updated columns.

## XLA Flags

For all update types, the `MapMessage` instance contains name/value pairs that indicate the following.

- Whether this is the first record of a transaction
- Whether this is the last record of a transaction
- Whether the update was performed by replication
- Which table was updated
- The owner of the updated table

The name/value pairs that contain these XLA flags are described in [Table 6-2](#). Each name is preceded by two underscores.

**Table 6-2 JMS/XLA Flags**

Name	Description	Corresponding <code>ttXlaUpdateDesc_t</code> Flag
<code>__AGING_DELETE</code>	Indicates that a delete was due to aging. The flag is present only if the XLA update record is due to an aging delete. The <code>XlaConstants</code> constant <code>AGING_DELETE_FIELD</code> represents this flag.	<code>TT_AGING</code>
<code>__CASCADING_DELETE</code>	Indicates that a delete was due to a cascading delete. The flag is present only if the XLA update record is due to a cascading delete. The <code>XlaConstants</code> constant <code>CASCADING_DELETE_FIELD</code> represents this flag.	<code>TT_CASCDEL</code>
<code>__COMMIT</code>	Indicates that this is the last record in a transaction and that a commit was performed after this operation. This is in the <code>MapMessage</code> if <code>TT_UPDCOMMIT</code> is on. The <code>XlaConstants</code> constant <code>COMMIT_FIELD</code> represents this flag.	<code>TT_UPDCOMMIT</code>
<code>__FIRST</code>	Indicates that this is the first record in a new transaction. This is in the <code>MapMessage</code> if <code>TT_UPDFIRST</code> is on. The <code>XlaConstants</code> constant <code>FIRST_FIELD</code> represents this flag.	<code>TT_UPDFIRST</code>
<code>__REPL</code>	Indicates that this change was applied to the database through replication. This is in the <code>MapMessage</code> if <code>TT_UPDREPL</code> is on. The <code>XlaConstants</code> constant <code>REPL_FIELD</code> represents this flag.	<code>TT_UPDREPL</code>

Table 6-2 (Cont.) JMS/XLA Flags

Name	Description	Corresponding ttXlaUpdateDesc_t Flag
__UPDCOLS	This is only used for UPDATETUP records, indicating that the XLA update descriptor contains a list of columns that were actually modified by the operation. It is specified as a string that contains a semicolon-delimited list of column names and is in the <code>MapMessage</code> only if <code>TT_UPDCOLS</code> is on. The <code>XlaConstants</code> constant <code>UPDATE_COLUMNS_FIELD</code> represents this flag.	<code>TT_UPDCOLS</code>

**Note**

The `XlaConstants` interface is in the `com.timesten.dataserver.jmsxla` package.

Applications can use the `MapMessage` method `itemExists()` to determine whether a flag is present, and `getBoolean()` to determine whether a flag is set. As input, specify the `XlaConstants` constant that corresponds to the flag, such as `XlaConstants.AGING_DELETE_FIELD`.

Equivalent to using `TT_UPDCOMMIT` in XLA, you can use the following test in JMS/XLA to see whether this is the last record in a transaction and that a commit was performed after the operation.

```
if (MapMessage.getBoolean(XlaConstants.COMMIT_FIELD) ) { // Field is set
    ...
}
```

## DML Event Data Formats

Many DML operations generate XLA updates that can be monitored by XLA event handlers.

This section describes the contents of the `MapMessage` instances that are generated for these operations.

### Table Data

For `INSERT`, `UPDATE` and `DELETE` operations, the `MapMessage` instance contains two name/value pairs, `__TBLOWNER` and `__TBLNAME`. These describe the name and owner of the table that is being updated.

For example, for a table `SCOTT.EMPLOYEES`, any related `MapMessage` contains a field `__TBLOWNER` with the string value `SCOTT` and a field `__TBLNAME` with the string value `EMPLOYEES`.

## Row Data

For `INSERT` and `DELETE` operations, a complete image of the inserted or deleted row is included in the message and all column values are available.

For `UPDATE` operations, the complete "before" and "after" images of the row are available, along with a list of column numbers indicating which columns were modified. Access the column values using the names of the columns. The column names in the "before" image all begin with a single underscore. For example, `columnname` contains the new value and `_columnname` contains the old value.

If the value of a column is `NULL`, it is omitted from the column list. The `__NULLS` name/value pair contains a semicolon-delimited list of the columns that contain `NULL` values.

## Context Information

If the `ttApplicationContext` built-in procedure was used to encode context information in an XLA record, that information is in the `__CONTEXT` name/value pair in the `MapMessage` instance. If no context information is provided, the `__CONTEXT` value is not in the `MapMessage` instance.

## DDL Event Data Formats

Many data definition language (DDL) operations generate XLA updates that can be monitored by XLA event handlers.

This section describes the contents of the `MapMessage` instances that are generated for these operations.

### CREATE\_TABLE

Messages with `__TYPE=1` (`XlaConstants.CREATE_TABLE`) indicate that a table has been created.

[Table 6-3](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `CREATE_TABLE` operation.

**Table 6-3 CREATE\_TABLE Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the created table
NAME	String value of the name of the created table
PK_COLUMNS	String value containing the names of the columns in the primary key for this table If the table has no primary key, the <code>PK_COLUMNS</code> value is not specified. Format: <code>&lt;col1name&gt;[;&lt;col2name&gt; [&lt;col3name&gt;[...]]]</code>

**Table 6-3 (Cont.) CREATE\_TABLE Data Provided in Update Messages**

Name	Value
COLUMNS	String value containing the names of the columns in the table Format:  <col1name>[ ; <col2name> [ ; <col3name> [ ; ... ] ]
	<b>Note:</b> For each column in the table, additional name/value pairs that describe the column are in the <code>MapMessage</code> .
<code>_column_name_TYPE</code>	Integer value representing the data type of this column (from <code>java.sql.Types</code> )
<code>_column_name_PRECISION</code>	Integer value containing the precision of this column (for NUMERIC or DECIMAL)
<code>_column_name_SCALE</code>	Integer value containing the scale of this column (for NUMERIC or DECIMAL)
<code>_column_name_SIZE</code>	Integer value indicating the maximum size of this column (for CHAR, VARCHAR, BINARY, or VARBINARY)
<code>_column_name_NULLABLE</code>	Boolean value indicating whether this column can have a NULL value
<code>_column_name_OUTOFFLINE</code>	Boolean value indicating whether this column is stored in the inline or out-of-line part of the tuple
<code>_column_name_INPRIMARYKEY</code>	Boolean value indicating whether this column is part of the primary key of the table

## DROP\_TABLE

Messages with `__TYPE=2` (`XlaConstants.DROP_TABLE`) indicate that a table has been dropped.

[Table 6-4](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `DROP_TABLE` operation.

**Table 6-4 DROP\_TABLE Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the sequence
NAME	String value of the name of the dropped sequence

## CREATE\_INDEX

Messages with `__TYPE=3` (`XlaConstants.CREATE_INDEX`) indicate that an index has been created.

[Table 6-5](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `CREATE_INDEX` operation.

**Table 6-5 CREATE\_INDEX Data Provided in Update Messages**

Name	Value
TBLOWNER	String value of the owner of the table on which the index was created
TBLNAME	String value of the name of the table on which the index was created
IXNAME	String value of the name of the created index
INDEX_TYPE	String value representing the index type: "P" (primary key), "F" (foreign key), or "R" (regular)
INDEX_METHOD	String value representing the index method: "H" (hash), "T" (range), or "B" (bit map)
UNIQUE	Boolean value indicating whether the index is unique
HASH_PAGES	Integer value representing the number of pages in a hash index (not specified for range indexes)
COLUMNS	String value describing the columns in the index Format: <code>&lt;col1name&gt;[ ;&lt;col2name&gt; [ ;&lt;col3name&gt;[ ;... ]]</code>

## DROP\_INDEX

Messages with `__TYPE=4` (`XlaConstants.DROP_INDEX`) indicate that an index has been dropped.

[Table 6-6](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `DROP_INDEX` operation.

**Table 6-6 DROP\_INDEX Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the table on which the index was dropped
TABLE_NAME	String value of the name of the table on which the index was dropped
INDEX_NAME	String value of the name of the dropped index

## ADD\_COLUMNS

Messages with `__TYPE=5` (`XlaConstants.ADD_COLUMNS`) indicate that a table has been altered by adding new columns.

[Table 6-7](#) shows the name/value pairs that are in a `MapMessage` instance generated for an `ADD_COLUMNS` operation.

**Table 6-7 ADD\_COLUMNS Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the altered table
NAME	String value of the name of the altered table

**Table 6-7 (Cont.) ADD\_COLUMNS Data Provided in Update Messages**

Name	Value
PK_COLUMNS	String value containing the names of the columns in the primary key for this table  If the table has no primary key, the PK_COLUMNS value is not specified.  Format:  <col1name>[;<col2name> [<col3name>[;...]]]
COLUMNS	String value containing the names of the columns added to the table  Format:  <col1name>[;<col2name> [<col3name>[;...]]]  <b>Note:</b> For each added column, additional name/value pairs that describe the column are in the MapMessage.
_column_name_TYPE	Integer value representing the data type of this column (from java.sql.Types)
_column_name_PRECISION	Integer value containing the precision of this column (for NUMERIC or DECIMAL)
_column_name_SCALE	Integer value containing the scale of this column (for NUMERIC or DECIMAL)
_column_name_SIZE	Integer value indicating the maximum size of this column (for CHAR, VARCHAR, BINARY, or VARBINARY)
_column_name_NULLABLE	Boolean value indicating whether this column can have a NULL value
_column_name_OUTOFFLINE	Boolean value indicating whether this column is stored in the inline or out-of-line part of the tuple
_column_name_INPRIMARYKEY	Boolean value indicating whether this column is part of the primary key of the table

## DROP\_COLUMNS

Messages with \_\_TYPE=6 (XlaConstants.DROP\_COLUMNS) indicate that a table has been altered by dropping existing columns.

[Table 6-8](#) shows the name/value pairs that are in a MapMessage instance generated for a DROP\_COLUMNS operation.

**Table 6-8 DROP\_COLUMNS Data Provided in Update Message**

Name	Value
OWNER	String value of the owner of the altered table
NAME	String value of the name of the altered table

**Table 6-8 (Cont.) DROP\_COLUMNS Data Provided in Update Message**

Name	Value
COLUMNS	String value containing the names of the columns dropped from the table Format: <code>&lt;col1name&gt;[ ;&lt;col2name&gt; [ ;&lt;col3name&gt;[ ; ... ]]</code> <b>Note:</b> For each dropped column, additional name/value pairs that describe the column are in the <code>MapMessage</code> .
<code>_column_name_TYPE</code>	Integer value representing the data type of this column (from <code>java.sql.Types</code> )
<code>_column_name_PRECISION</code>	Integer value containing the precision of this column (for NUMERIC or DECIMAL)
<code>_column_name_SCALE</code>	Integer value containing the scale of this column (for NUMERIC or DECIMAL)
<code>_column_name_SIZE</code>	Integer value indicating the maximum size of this column (for CHAR, VARCHAR, BINARY, or VARBINARY)
<code>_column_name_NULLABLE</code>	Boolean value indicating whether this column can have a NULL value
<code>_column_name_OUTOFFLINE</code>	Boolean value indicating whether this column is stored in the inline or out-of-line part of the tuple
<code>_column_name_INPRIMARYKEY</code>	Boolean value indicating whether this column is part of the primary key of the table

## CREATE\_VIEW

Messages with `__TYPE=14` (`XlaConstants.CREATE_VIEW`) indicate that a materialized view has been created.

[Table 6-9](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `CREATE_VIEW` operation.

**Table 6-9 CREATE\_VIEW Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the created view
NAME	String value of the name of the created view

## DROP\_VIEW

Messages with `__TYPE=15` (`XlaConstants.DROP_VIEW`) indicate that a materialized view has been dropped.

See [Table 6-10](#) to see the name/value pairs that are in a `MapMessage` instance generated for a `DROP_VIEW` operation.

**Table 6-10 DROP\_VIEW Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the dropped view
NAME	String value of the name of the dropped view

## CREATE\_SEQ

Messages with `__TYPE=16` (`XlaConstants.CREATE_SEQ`) indicate that a sequence has been created.

[Table 6-11](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `CREATE_SEQ` operation.

**Table 6-11 CREATE\_SEQ Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the created sequence
NAME	String value of the name of the created sequence
CYCLE	Boolean value indicating whether the <code>CYCLE</code> option was specified on the new sequence
INCREMENT	A long value indicating the <code>INCREMENT BY</code> option specified for the new sequence
MIN_VALUE	A long value indicating the <code>MINVALUE</code> option specified for the new sequence
MAX_VALUE	A long value indicating the <code>MAXVALUE</code> option specified for the new sequence

## DROP\_SEQ

Messages with `__TYPE=17` (`XlaConstants.DROP_SEQ`) indicate that a sequence has been dropped.

[Table 6-12](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `DROP_SEQ` operation.

**Table 6-12 DROP\_SEQ Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the dropped table
NAME	String value of the name of the dropped table

## CREATE\_SYNONYM

Messages with `__TYPE=19` (`XlaConstants.CREATE_SYNONYM`) indicate that a synonym has been created.

[Table 6-13](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `CREATE_SYNONYM` operation.

**Table 6-13 CREATE\_SYNONYM Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the created synonym
NAME	String value of the name of the created synonym
OBJECT_OWNER	String value of the schema of the object for which you are creating a synonym
OBJECT_NAME	String value of the name of the object for which you are creating a synonym
IS_PUBLIC	Boolean value indicating whether the synonym is public
IS_REPLACE	Boolean value indicating whether the synonym was created using CREATE OR REPLACE

## DROP\_SYNONYM

Messages with `__TYPE=20` (`XlaConstants.DROP_SYNONYM`) indicate that a synonym has been dropped.

[Table 6-14](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `DROP_SYNONYM` operation.

**Table 6-14 DROP\_SYNONYM Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the dropped synonym
NAME	String value of the name of the dropped synonym
IS_PUBLIC	Boolean value indicating whether the synonym was public

## TRUNCATE

Messages with `__TYPE=18` (`XlaConstants.TRUNCATE`) indicate that a table has been truncated. All rows in the table have been deleted.

[Table 6-15](#) shows the name/value pairs that are in a `MapMessage` instance generated for a `TRUNCATE` operation.

**Table 6-15 TRUNCATE Data Provided in Update Messages**

Name	Value
OWNER	String value of the owner of the truncated table
NAME	String value of the name of the truncated table

## Data Type Support

This section covers data type considerations for JMS/XLA.

## Data Type Mapping

There are `MapMessage` access methods for data types supported by TimesTen.

See [Table 6-16](#). For more information about data types, see Data Types in *Oracle TimesTen In-Memory Database SQL Reference*.

**Table 6-16 Data Type Mapping**

TimesTen column type	Read with MapMessage method...
CHAR( <i>n</i> )	<code>getString()</code>
VARCHAR( <i>n</i> )	<code>getString()</code>
NCHAR( <i>n</i> )	<code>getString()</code>
NVARCHAR( <i>n</i> )	<code>getString()</code>
NVARCHAR2( <i>n</i> )	<code>getString()</code>
DOUBLE	<code>getString()</code> Can be converted to <code>BigDecimal</code> or to <code>Double</code> by the application.
FLOAT	<code>getString()</code> Can be converted to <code>BigDecimal</code> or to <code>Double</code> by the application.
DECIMAL( <i>p,s</i> )	<code>getString()</code> Can be converted to <code>BigDecimal</code> or to <code>Double</code> by the application.
NUMERIC( <i>p,s</i> )	<code>getString()</code> Can be converted to <code>BigDecimal</code> or to <code>Double</code> by the application.
INTEGER	<code>getInt()</code>
SMALLINT	<code>getShort()</code>
TINYINT	<code>getShort()</code>
BINARY( <i>n</i> )	<code>getBytes()</code>
VARBINARY( <i>n</i> )	<code>getBytes()</code>
DATE	<code>getLong()</code> , <code>getString()</code> The <code>getLong()</code> method returns microseconds since epoch (00:00:00 UTC, January 1, 1970). Can be converted to <code>Date</code> or <code>Calendar</code> by the application.
TIME	<code>getString()</code> Can be converted to <code>Date</code> or <code>Calendar</code> by the application.
TIMESTAMP	<code>getLong()</code> , <code>getString()</code> The <code>getLong()</code> method returns microseconds since epoch (00:00:00 UTC, January 1, 1970). It truncates nanoseconds. Use <code>getString()</code> if you require nanosecond precision. Can be converted to <code>Date</code> or <code>Calendar</code> by the application.
TT_CHAR	<code>getString()</code>
TT_VARCHAR	<code>getString()</code>
TT_NCHAR	<code>getString()</code>
TT_NVARCHAR	<code>getString()</code>
ORA_CHAR	<code>getString()</code>

**Table 6-16 (Cont.) Data Type Mapping**

TimesTen column type	Read with MapMessage method...
ORA_VARCHAR	getString()
ORA_NCHAR	getString()
ORA_NVARCHAR2	getString()
VARCHAR2	getString()
TT_TINYINT	getShort()
TT_SMALLINT	getShort()
TT_INTEGER	getInt()
TT_BIGINT	getLong()
BINARY_FLOAT	getFloat()
BINARY_DOUBLE	getDouble()
REAL	getFloat()
NUMBER	getString()
ORA_NUMBER	getString()
TT_TIME	getString()
TT_DATE	getLong(), getString() The getLong() method returns microseconds since epoch (00:00:00 UTC, January 1, 1970).
TT_TIMESTAMP	getLong(), getString() The getLong() method returns microseconds since epoch (00:00:00 UTC, January 1, 1970).
ORA_DATE	getLong(), getString() The getLong() method returns microseconds since epoch (00:00:00 UTC, January 1, 1970).
ORA_TIMESTAMP	getLong(), getString() The getLong() method returns microseconds since epoch (00:00:00 UTC, January 1, 1970). It truncates nanoseconds. Use getString() if you require nanosecond precision.
TT_BINARY	getBytes()
TT_VARBINARY	getBytes()
ROWID	getBytes(), getString()
BLOB	getBytes() <b>Note:</b> Information about the LOB value itself is unavailable. LOB fields contain zero-length data or null data (if the value is actually NULL).
CLOB, NCLOB	getString() <b>Note:</b> Information about the LOB value itself is unavailable. LOB fields contain zero-length data or null data (if the value is actually NULL).

## Data Types Character Set

JMS/XLA uses a UTF-16 character set for specific data types.

- TT\_CHAR
- TT\_VARCHAR
- ORA\_CHAR
- ORA\_VARCHAR2
- TT\_NCHAR
- TT\_NVARCHAR
- ORA\_NCHAR
- ORA\_NVARCHAR2
- NCHAR
- NVARCHAR
- NVARCHAR2

## JMS Interfaces for Event Handling

The standard JMS interfaces are available for JMS/XLA applications. Note that the JMS/XLA API supports only publish/subscribe messaging.

- Connection
- ConnectionFactory
- ConnectionMetaData
- Destination
- ExceptionListener
- MapMessage
- Message
- MessageConsumer
- Session
- Topic
- TopicConnection
- TopicConnectionFactory
- TopicSession
- TopicSubscriber

See the following API documentation for these interfaces:

- <https://javadoc.io/doc/jakarta.platform/jakarta.jakartaee-api/9.0.0/index.html> to review the `jakarta.jms` package documentation for the Jakarta Messaging API version 1.1 for Jakarta JMS. Note TimesTen doesn't utilize the other Jakarta Platform EE 9 APIs.
- <https://javaee.github.io/javaee-spec/javadocs/> to review the `javax.jms` package documentation for the Java Message Service API version 1.1 for JavaX JMS

## JMS/XLA Replication API

The TimesTen `com.timesten.dataserver.jmsxla` package includes the `TargetDataStore` interface and the `TargetDataStoreImpl` class, described in this section.

See *Oracle TimesTen In-Memory Database JavaX JMS/XLA Java API Reference*.

### TargetDataStore Interface

This interface is used to apply XLA update records from a source database to a target database. The source and target database schema must be identical for the affected tables.

This interface defines the methods shown in [Table 6-17](#).

**Table 6-17 TargetDataStore Methods**

Method	Description
<code>apply()</code>	Applies XLA update descriptor to the target database.
<code>close()</code>	Closes the connections to the database and releases the resources.
<code>commit()</code>	Performs a manual commit.
<code>getAutoCommitFlag()</code>	Returns the value of the autocommit flag.
<code>getConnectString()</code>	Returns the database connection string.
<code>getUpdateConflictCheckFlag()</code>	Returns the value of the flag for checking update conflicts.
<code>isClosed()</code>	Checks whether the object is closed.
<code>isDataStoreValid()</code>	Checks whether the database is valid.
<code>rollback()</code>	Rolls back the last transaction.
<code>setAutoCommitFlag()</code>	Sets the flag for autocommit during apply.
<code>setUpdateConflictCheckFlag()</code>	Sets the flag for checking update conflicts during apply.

### TargetDataStoreImpl Class

This class creates connections and XLA handles for a target database. It implements the `TargetDataStore` interface.

## JMS Message Header Fields

Description of JMS messages header fields provided by JMS/XLA.

**Table 6-18 JMS/XLA Header Fields**

Header	Contents
<code>JMSMessageId</code>	Transaction log file number of the XLA record
<code>JMSType</code>	String representation of the <code>__TYPE</code> field