

*Oracle TimesTen
In-Memory Database
Operations Guide*

Release 7.0

B31689-03



Copyright ©1996, 2007, Oracle. All rights reserved.

ALL SOFTWARE AND DOCUMENTATION (WHETHER IN HARD COPY OR ELECTRONIC FORM) ENCLOSED AND ON THE COMPACT DISC(S) ARE SUBJECT TO THE LICENSE AGREEMENT.

The documentation stored on the compact disc(s) may be printed by licensee for licensee's internal use only. Except for the foregoing, no part of this documentation (whether in hard copy or electronic form) may be reproduced or transmitted in any form by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without the prior written permission of TimesTen Inc.

Oracle, JD Edwards, PeopleSoft, Retek, TimesTen, the TimesTen icon, MicroLogging and Direct Data Access are trademarks or registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

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

September 2007

Printed in the United States of America

Contents

About this Guide

TimesTen documentation	1
Background reading	3
Installing TimesTen	3
Conventions used in this guide	4
Finding the information you need	5
Technical Support	7

1 Creating TimesTen Data Stores

TimesTen ODBC and JDBC drivers	10
TimesTen ODBC drivers	10
TimesTen JDBC driver	12
JDBC driver manager	12
Data source names	13
User and system DSNs	13
Data Manager and Client DSNs	14
Connection attributes for Data Manager DSNs	15
Thread programming with TimesTen	16
Creating a DSN on Windows	17
Specify the ODBC driver	17
Specify the DSN	17
Specify the connection attributes	18
Creating a DSN on UNIX	22
Create a user ODBC.INI file	22
Specify the DSN	22
Specify the ODBC driver	22
Specify the data store path name	23
Choose a database character set	23
Set data store attributes	23
Using environment variables in data store path names	23
DSN examples	24
Setting up a temporary data store	24
Creating multiple DSNs to a single data store	25
Connecting to a data store without using a DSN	28
Specifying the size of a data store	29
Temporary and permanent memory	29
Changing data store size	29
Receiving out-of-memory warnings.	30

Specifying a RAM policy31
Copying, migrating, backing up and restoring a data store32
Working with the ODBC.INI file34
The user ODBC.INI file35
The system ODBC.INI file35
Searching for a DSN35
ODBC Data Sources35
Data Source Specification35
odbc.ini file example37

2 Working with the TimesTen Client and Server

Client/Server Communication40
TCP/IP Communication40
Shared memory communication40
UNIX domain socket communication40
Configuring TimesTen Client and Server41
Configuring Client/Server of the same TimesTen release.41
Configuring cross-release Client/Server42
Running the TimesTen Server Daemon42
Server informational messages43
TimesTen Server connection attributes43
Defining Server DSNs43
TimesTen Client connection attributes44
Creating and configuring Client DSNs on Windows45
Creating and configuring a logical server name45
Creating a Client DSN on Windows46
Setting the timeout interval and authentication.48
Deleting a server name48
Accessing a remote data store on Windows49
Testing connections50
Creating and configuring Client DSNs on UNIX51
Searching for a DSN52
Creating and configuring a logical server name52
Examples53
Creating a Client DSN54
Accessing a remote data store on UNIX.54
Testing connections56
Working with the TTCONNECT.INI file57
Defining a server name on UNIX57

3 Working with the Oracle TimesTen Data Manager Daemon

Starting and stopping the Oracle TimesTen Data Manager	
Service on Windows60
Starting and stopping the daemon on UNIX60
Managing TimesTen daemon options61
Determining the daemon listening address62
Listening on IPv663
Modifying informational messages64
Changing the allowable number of subdaemons66
Allowing data store access over NFS-mounted systems66
Enabling Linux large page support66
Other daemon options settings66
Managing TimesTen Client/Server daemon options67
Modifying the TimesTen Server daemon options67
Controlling the TimesTen Server daemon.67
Prespawning TimesTen Server processes68
Specifying multiple connections to the TimesTen Server.68
Configuring the maximum number of client connections per child server process68
Configuring the desired number of child server processes spawned for a server DSN69
Configuring the Thread Stack Size of the Child Server Processes69
Using shared memory for Client/Server IPC69
Managing the size of the shared memory segment70
Changing the size of the shared memory segment71
Controlling the TimesTen Server log messages71
Communicating with older versions of TimesTen71
Modifying the TimesTen web server options72
Controlling the TimesTen web server72

4 Globalization Support

Overview of globalization support features75
Choosing a database character set76
Character sets and languages.76
Client operating system and application compatibility.77
Performance and storage implications77
Character sets and replication77
Length semantics and data storage77
Connection character set.78
Linguistic sorts79
Monolingual linguistic sorts79
Multilingual linguistic sorts80

Case-insensitive and accent-insensitive linguistic sorts80
Performing a linguistic sort81
Using linguistic indexes81
SQL string and character functions.82
Setting globalization support attributes83
Backward compatibility using TIMESTEN884
Globalization support during migration84
Object migration and character sets84
Migration and length semantics85
Migrating linguistic indexes85
Migrating cache group tables85
Supported Character Sets86
Asian Character Sets86
European Character Sets86
Middle Eastern Character Sets88
TimesTen Character Set88
Universal Character Sets89
Supported Linguistic Sorts90
Monolingual Linguistic Sorts90
Multilingual Linguistic Sorts92

5 Using the ttIsql Utility

Batch mode vs. interactive mode97
Customizing the ttIsql command prompt99
Using ttIsql's online help	100
Using ttIsql's 'editline' feature (UNIX only)	102
Emacs binding	102
vi binding	103
Using ttIsql's command history	104
Saving and clearing ttIsql's command history	105
Working with character sets	106
Working with transactions	107
Displaying data store information	109
Viewing and setting data store attributes	112
Viewing and changing query optimizer plans.	113
Timing ODBC function calls	117
Working with prepared and parameterized SQL statements	118
Defining default settings with the TTISQL environment variable	122
Managing XLA bookmarks	124

6 Working with Data in a TimesTen Data Store

Data store overview	125
-------------------------------	-----

Data store components	126
Data store users and owners	126
Data store persistence	126
Understanding tables	127
In-line and out-of-line columns.	127
Default column values	128
Table names	128
Table access	129
Primary keys, foreign keys and unique indexes	129
System tables.	129
Working with tables	130
Creating a table	130
Dropping a table	130
Estimating table size	130
Implementing aging in your tables	131
Usage-based aging.	131
Time-based aging	133
Aging and foreign keys	134
Scheduling when aging starts	134
Aging and replication	135
Understanding materialized views	135
Working with materialized views	136
Creating a materialized view.	136
The SELECT query in the CREATE MATERIALIZED VIEW statement	137
Restrictions on materialized views and detail tables.	138
Performance implications of materialized views	139
Dropping a materialized view	140
Understanding views	140
Working with views	141
Creating a view	141
The SELECT query in the CREATE VIEW statement.	141
Restrictions on views and their detail tables.	142
Dropping a view	142
Understanding indexes	142
Working with indexes	143
Creating an index	144
Altering an index	144
Dropping an index.	144
Estimating index size	144
Understanding rows	144

Working with rows	145
Inserting rows	145
Deleting rows	145

7 Transaction Management and Recovery

Transaction semantics	148
Transaction atomicity and durability	151
Overview	151
Guaranteed atomicity and durability	152
Guaranteed atomicity, delayed durability	153
No guaranteed atomicity, no guaranteed durability	153
Controlling durability and logging	155
Using durable commits	155
Log files	156
Concurrency control	157
Transaction isolation levels	157
Locking granularities	159
Coexistence of different locking levels	159
Checkpoints	160
Types of checkpoints	161
Transaction-consistent checkpoints	161
Fuzzy or non-blocking checkpoints	161
Setting and managing checkpoints	162
Setting the checkpoint rate for background checkpoints	163

8 Data Store Performance Tuning

System and data store tuning	166
Provide enough memory	166
Size your data store correctly	166
Increase LogBuffSize if needed	166
Use temporary data stores if appropriate	167
Avoid connection overhead	167
Load the data store into RAM when duplicating	167
Avoid OS paging at load time	168
Consider special options for maintenance	168
Check your driver	168
Enable tracing only as needed	169
Investigate alternative JVMs	169
Adjust log buffer size and CPU for a large number of subscribers	169
Migrating data with character set conversions	169
Client/Server tuning	170
Work locally when possible	170

Use shared memory segment as IPC when client and server are on the same machine	170
Enable TT_PREFETCH_CLOSE for serializable transactions	171
Use a connection handle when calling SQLTransact	173
SQL tuning	174
Tune statements and use indexes	174
Select hash or T-tree indexes appropriately	175
Size hash indexes appropriately	176
Use foreign key constraint appropriately	176
Computing exact or estimated statistics	176
Avoid ALTER TABLE	177
Avoid nested queries	177
Improving performance of materialized views	178
Limit number of join rows	178
Use indexes on join columns.	178
Avoid unnecessary updates	179
Avoid changes to the inner table of an outer join	179
Limit number of columns in a view table	179
Scaling to Multiple CPUs	181
Run the demo applications as a prototype.	181
Limit database-intensive connections per CPU	181
Use read operations when available	182
Limit prepares, re-prepares and connects	182
Limit replication transmitters and receivers and XLA readers	182
Allow indexes to be rebuilt in parallel during recovery	183
Use private commands	183
XLA acknowledgement modes	184
Prefetch multiple update records	184
Acknowledge XLA updates	184

9 The TimesTen Query Optimizer

When optimization occurs	186
Viewing a plan	189
Generating the plan	189
Reading the PLAN table	189
PLAN table columns.	191
Modifying plan generation	193
Why modify an execution plan?	193
When to modify an execution plan	193
How to modify execution plan generation	197

Glossary

Index

About this Guide

Oracle TimesTen In-Memory Database is a high-performance, in-memory data manager that supports the ODBC and JDBC interfaces.

This guide provides:

- Background information to help you understand how TimesTen works.
- Step-by-step instruction and examples that show how to perform the most commonly needed tasks.

To work with this guide, you should understand how database systems work and have some knowledge of SQL (Structured Query Language).

TimesTen documentation

TimesTen documentation is available on the product distribution media and on the Oracle Technology Network:

http://www.oracle.com/technology/documentation/timesten_doc.html.

Including this guide, the TimesTen documentation set consists of these documents:

Book Titles	Description
<i>Oracle TimesTen In-Memory Database Installation Guide</i>	Contains information needed to install and configure TimesTen on all supported platforms.
<i>Oracle TimesTen In-Memory Database Introduction</i>	Describes all the available features in the Oracle TimesTen In-Memory Database.
<i>Oracle TimesTen In-Memory Database Operations Guide</i>	Provides information on configuring TimesTen and using the ttlsq utility to manage a data store. This guide also provides a basic tutorial for TimesTen.
<i>Oracle TimesTen In-Memory Database C Developer's and Reference Guide</i> and the <i>Oracle TimesTen In-Memory Database Java Developer's and Reference Guide</i>	Provide information on how to use the full set of available features in TimesTen to develop and implement applications that use TimesTen.
<i>Oracle TimesTen In-Memory Database API Reference Guide</i>	Describes all TimesTen utilities, procedures, APIs and provides a reference to other features of TimesTen.

<i>Oracle TimesTen In-Memory Database SQL Reference Guide</i>	Contains a complete reference to all TimesTen SQL statements, expressions and functions, including TimesTen SQL extensions.
<i>Oracle TimesTen In-Memory Database Error Messages and SNMP Traps</i>	Contains a complete reference to the TimesTen error messages and information on using SNMP Traps with TimesTen.
<i>Oracle TimesTen In-Memory Database TTClasses Guide</i>	Describes how to use the TTClasses C++ API to use the features available in TimesTen to develop and implement applications.
<i>TimesTen to TimesTen Replication Guide</i>	Provides information to help you understand how TimesTen Replication works and step-by-step instructions and examples that show how to perform the most commonly needed tasks. This guide is for application developers who use and administer TimesTen and for system administrators who configure and manage TimesTen Replication.
<i>TimesTen Cache Connect to Oracle Guide</i>	Describes how to use Cache Connect to cache Oracle data in TimesTen data stores. This guide is for developers who use and administer TimesTen for caching Oracle data.
<i>Oracle TimesTen In-Memory Database Troubleshooting Procedures Guide</i>	Provides information and solutions for handling problems that may arise while developing applications that work with TimesTen, or while configuring or managing TimesTen.

Background reading

For a Java reference, see:

- Horstmann, Cay and Gary Cornell. *Core Java(TM) 2, Volume I-- Fundamentals (7th Edition) (Core Java 2)*. Prentice Hall PTR; 7 edition (August 17, 2004).

A list of books about ODBC and SQL is in the Microsoft ODBC manual included in your developer's kit. Your developer's kit includes the appropriate ODBC manual for your platform:



- *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide* provides all relevant information on ODBC for Windows developers.
- *Microsoft ODBC 2.0 Programmer's Reference and SDK Guide*, included online in PDF format, provides information on ODBC for UNIX developers.

For a conceptual overview and programming how-to of ODBC, see:

- Kyle Geiger. *Inside ODBC*. Redmond, WA: Microsoft Press. 1995.

For a review of SQL, see:

- Melton, Jim and Simon, Alan R. *Understanding the New SQL: A Complete Guide*. San Francisco, CA: Morgan Kaufmann Publishers. 1993.
- Groff, James R. / Weinberg, Paul N. *SQL: The Complete Reference, Second Edition*. McGraw-Hill Osborne Media. 2002.

For information about Unicode, see:

- The Unicode Consortium, *The Unicode Standard, Version 5.0*, Addison-Wesley Professional, 2006.
- The Unicode Consortium Home Page at <http://www.unicode.org>

Installing TimesTen

TimesTen Release 7.0 includes the TimesTen Data Manager for 32-bit and 64-bit platforms. See the [Oracle TimesTen In-Memory Database Installation Guide](#) for a description of supported platforms.

In addition to the Data Manager, TimesTen Release 7.0 also includes TimesTen Client and Server components. You can install the TimesTen Data Manager stand-alone or in a client/server environment.

For a list of the The TimesTen default installation directories, see the [Oracle TimesTen In-Memory Database Installation Guide](#).

Conventions used in this guide

TimesTen supports multiple platforms. Unless otherwise indicated, the information in this guide applies to all supported platforms. The term Windows refers to Windows 2000, Windows XP and Windows Server 2003. The term UNIX refers to Solaris, Linux, HP-UX, Tru64 and AIX.

TimesTen documentation uses these typographical conventions:

If you see...	It means...
<code>code font</code>	Code examples, filenames, and pathnames. For example, the <code>.odbc.ini</code> or <code>ttconnect.ini</code> file.
<i>italic code font</i>	A variable in a code example that you must replace. For example: <code>Driver=install_dir/lib/libtten.sl</code> Replace <i>install_dir</i> with the path of your TimesTen installation directory.

TimesTen documentation uses these conventions in command line examples and descriptions:

If you see...	It means...
<i>fixed width italics</i>	Variable; must be replaced with an appropriate value. In some cases, such as for parameter values in built-in procedures, you may need to single quote (' ') the value.
[]	Square brackets indicate that an item in a command line is optional.
{ }	Curly braces indicated that you must choose one of the items separated by a vertical bar () in a command line.
	A vertical bar (or pipe) separates arguments that you may use more than one argument on a single command line.
...	An ellipsis (. . .) after an argument indicates that you may use more than one argument on a single command line.
%	The percent sign indicates the UNIX shell prompt.
#	The number (or pound) sign indicates the UNIX root prompt.

TimesTen documentation uses these variables to identify path, file and user names:

If you see...	It means...
<i>install_dir</i>	The path that represents the directory where the current release of TimesTen is installed.
<i>TTinstance</i>	The instance name for your specific installation of TimesTen. Each installation of TimesTen must be identified at install time with a unique alphanumeric instance name. This name appears in the install path. The instance name “giraffe” is used in examples in this guide.
<i>bits</i> or <i>bb</i>	Two digits, either 32 or 64, that represent either the 32-bit or 64-bit operating system.
<i>release</i> or <i>rr</i>	Two digits that represent the first two digits of the current TimesTen release number, with or without a dot. For example, 70 or 7.0 represents TimesTen Release 7.0.
<i>jdk_version</i>	Two digits that represent the version number of the major JDK release. Specifically, 14 represent JDK 1.4; 5 represents JDK 5.
<i>timesten</i>	A sample name for the TimesTen instance administrator. You can use any legal user name as the TimesTen administrator. On Windows, the TimesTen instance administrator must be a member of the Administrators group. Each TimesTen instance can have a unique instance administrator name.
<i>DSN</i>	The data source name.

Finding the information you need

The table below provides a brief overview of the TimesTen development process. It will help you get started with your application and find relevant information as you progress.

To learn how to	See
Install TimesTen	Oracle TimesTen In-Memory Database Installation Guide
Create and use a TimesTen data store	Chapter 1, “QuickStart”

To learn how to	See
Use the ttIsql command-line utility to execute SQL	Chapter 5, “Using the ttIsql Utility”
Manage TimesTen resources	Chapter 3, “Working with the Oracle TimesTen Data Manager Daemon”
Troubleshoot problems running the TimesTen demos	Oracle TimesTen In-Memory Database Troubleshooting Procedures Guide

Technical Support

For information about obtaining technical support for TimesTen products, go to the following Web address:

<http://www.oracle.com/support/contact.html>

Creating TimesTen Data Stores

A TimesTen data store is a collection of tables and indexes that can be accessed and manipulated through SQL. This chapter describes how to set up a TimesTen data store. It begins with an overview of the things you should consider when setting up a data store and then describes each task in detail.

Once you have created a data store, you can:

- Use the **ttIsql** utility to connect to the data store and execute a SQL file or start an interactive SQL session, as described in “[Batch mode vs. interactive mode](#)” on page 97.
- Execute an application that uses the data store, as described *Oracle TimesTen In-Memory Database Java Developer’s and Reference Guide* and *Oracle TimesTen In-Memory Database C Developer’s and Reference Guide*.

The main topics are:

- [TimesTen ODBC and JDBC drivers](#)
- [Data source names](#)
- [Creating a DSN on Windows](#)
- [Creating a DSN on UNIX](#)
- [DSN examples](#)
- [Specifying the size of a data store](#)
- [Specifying a RAM policy](#)
- [Copying, migrating, backing up and restoring a data store](#)
- [Working with the ODBC.INI file](#)

TimesTen ODBC and JDBC drivers

This section describes some basic concepts that will help you define TimesTen data stores. The main topics are:

- [TimesTen ODBC drivers](#)
- [TimesTen JDBC driver](#)

C applications interact with TimesTen either by linking directly with a TimesTen ODBC driver, or by linking with an ODBC driver manager. Java applications access the ODBC driver through a JDBC library. Consider the following points:

- An application that uses an ODBC driver manager dynamically loads an ODBC driver at runtime. A single application can use more than one ODBC driver within the same application, even if the drivers are for different RDBMS products. The downside to this flexibility is that the driver manager adds additional runtime overhead.
- An application that links directly with an ODBC driver can use only the driver with which it is linked. This option offers less flexibility but better performance than linking with a driver manager.
- An application that is using an ODBC driver manager cannot use XLA.



On Windows, TimesTen installs the ODBC driver manager if it is not already present on the machine.



On UNIX, ODBC driver managers are available from third-party vendors. ODBC driver managers are not supplied with UNIX systems. This guide refers to the use of a driver manager for Windows systems only.

For more information on how to compile an application that uses the TimesTen data manager, see [Oracle TimesTen In-Memory Database Java Developer's and Reference Guide](#) and [Oracle TimesTen In-Memory Database C Developer's and Reference Guide](#).

TimesTen ODBC drivers

TimesTen includes two versions of the Data Manager ODBC driver: a *production* version and a *debug* version.

- Use the *production* version of the TimesTen Data Manager driver for most application development and for all deployment.
- Use the *debug* version of the TimesTen Data Manager driver only if you encounter problems with TimesTen itself. This version performs additional internal error checking and is slower than the production version. (On UNIX, the TimesTen debug libraries are compiled with the `-g` option to display additional debug information.)

TimesTen also includes the TimesTen Client ODBC driver for use with client/server applications.



On Windows, the production version of the TimesTen Data Manager is installed by default. To install the debug version, choose **Custom** setup. To install the TimesTen Client driver, choose either **Typical** or **Custom** setup. The following table lists the ODBC drivers for Windows:

Platform	Version	Name
Windows	Production	TimesTen Data Manager 7.0 Driver.
Windows	Debug	TimesTen Data Manager 7.0 Debug Driver.
Windows	Client	TimesTen Client 7.0 Driver



On UNIX, depending on the options selected at install time, TimesTen installs the Client driver and/or both the production version and the debug version of the TimesTen Data Manager ODBC driver. The following table lists the TimesTen ODBC drivers for UNIX platforms, assuming TimesTen is installed in the default directory by user root.

Platform	Version	Location and Name
HP-UX	Production	<code>/opt/TimesTen/TTinstance/lib/libtten.sl</code> TimesTen Data Manager 7.0 Driver.
HP-UX	Debug	<code>/opt/TimesTen/TTinstance/lib/libttenD.sl</code> TimesTen Data Manager 7.0 Debug Driver.
HP-UX	Client	<code>/opt/TimesTen/TTinstance/lib/libttclient.sl</code> TimesTen Client 7.0 Driver.
Solaris Linux Tru64	Production	<code>/opt/TimesTen/TTinstance/lib/libtten.so</code> TimesTen Data Manager 7.0 Driver.
Solaris Linux Tru64	Debug	<code>/opt/TimesTen/TTinstance/lib/libttenD.so</code> TimesTen Data Manager 7.0 Debug Driver.
Solaris Linux Tru64	Client	<code>/opt/TimesTen/TTinstance/lib/libttclient.so</code> TimesTen Client 7.0 Driver.
AIX	Production	<code>/usr/lpp/TimesTen/TTinstance/lib/libtten.a</code> TimesTen Data Manager 7.0 Driver.

Platform	Version	Location and Name
AIX	Debug	<code>/usr/lpp/TimesTen/TTinstance/lib/libttenD.a</code> TimesTen Data Manager 7.0 Debug Driver.
AIX	Client	<code>/usr/lpp/TimesTen/TTinstance/lib/libttclient.a</code> TimesTen Client 7.0 Driver.

TimesTen JDBC driver

The TimesTen JDBC driver uses the ODBC driver to access the TimesTen data stores. For each JDBC method, the driver executes a set of ODBC functions to perform the appropriate operation. Since the JDBC driver depends on ODBC for all data store operations, the first step in using JDBC is to define a TimesTen data store and the ODBC driver that will access it on behalf of JDBC.

JDBC is installed with the TimesTen Data Manager. JDBC allows Java applications to issue SQL statements to TimesTen and process the results. JDBC is the primary interface for data access in the Java programming language.

The JDBC API is implemented using a driver manager that can support multiple drivers connecting to different databases. The TimesTen JDBC driver is implemented using native methods to bridge to the TimesTen native API.

For a list of the functions supported by TimesTen, see the [Oracle TimesTen In-Memory Database Java Developer's and Reference Guide](#).



JDBC driver manager

The JDBC driver manager (DriverManager class) keeps track of all the JDBC drivers that have been loaded and are available to the Java Application. The application may load several drivers and access each driver independently. For example, both the TimesTen JDBC Client/Server driver and the TimesTen JDBC direct driver can be loaded onto a machine. Then Java applications can access data stores either on the local machine or a remote machine.

Data source names

TimesTen data stores are accessed through Data Source Names (DSNs). A DSN is a character-string name that identifies a TimesTen data store and a collection of connection attributes that are to be used when connecting to the data store. On Windows, the DSN also specifies the ODBC driver to be used to access the data store.

Each DSN uniquely identifies a data store. However, a data store can be referenced by multiple DSNs. Each of these DSNs can specify a different set of connection attributes. This allows you to give convenient names to different connection configurations for a single data store.

Note: According to the ODBC standard, when an attribute occurs multiple times in a connection string, the first value specified is used, not the last value.

A DSN has the following characteristics:

- Its maximum length is 32 characters.
- It is composed of ASCII characters except for the following: []{};,*=!@\
- It cannot contain spaces.

The rest of this section includes the following topics:

- [User and system DSNs](#)
- [Data Manager and Client DSNs](#)
- [Connection attributes for Data Manager DSNs](#)
- [Thread programming with TimesTen](#)

User and system DSNs

DSNs are resolved using a two-tiered naming system, consisting of user DSNs and system DSNs:

- A **user DSN** can be used only by the user who created the DSN. On Windows, user DSNs are defined from the **User DSN** tab of the ODBC Data Source Administrator. On UNIX, user DSNs are defined in the file `$HOME/.odbc.ini` or in a file named by the ODBCINI environment variable. This file is referred to as the “user ODBC.INI file.” Although a user DSN is private to the user who created it, it is only the DSN (the character-string name and its attributes) that is private. The underlying data store can be referenced by other users’ user DSNs or by system DSNs.
- A **system DSN** can be used by any user on the machine on which the system DSN is defined. On Windows, system DSNs are defined from the **System DSN** tab of the ODBC Data Source Administrator. On UNIX, system DSNs are defined in the file `/var/TimesTen/sys.odbc.ini` if TimesTen is installed

by user `root`, or `install_dir/info/sys.odbc` if TimesTen is installed by a non-root user. This file is referred to as the “system ODBC.INI file.”

When looking for a specific DSN, TimesTen first looks for a user DSN with the specified name. If no matching user DSN is found, TimesTen looks for a system DSN with the specified name. If a user DSN and a system DSN with the same name exist, TimesTen uses the user DSN. On UNIX, if there are multiple DSNs with the same name in the same ODBC.INI file, TimesTen uses the first one in the file.

Data Manager and Client DSNs

DSNs that use the TimesTen Data Manager (either the production version or the debug version) are called “Data Manager DSNs.” DSNs that use the TimesTen Client are called “Client DSNs.”

Data Manager DSNs define what is referred to as a *direct driver* connection to the data store. A Data Manager DSN refers to a data store using a path name. A data store path name is a path name that specifies the location of the data store, for example: `C:\data\chns\AdminDS` or `/home/chns/AdminDS`. This name is not a file name. The actual files used by the data store have file suffixes, for example: `C:\data\chns\AdminDS.ds0` or `/home/chns/AdminDS.log2`. A Data Manager DSN that refers to a given TimesTen data store must be defined on the same system on which the data store resides. In addition, TimesTen creates `dsName.resn` files for each data store. These files are used internally by TimesTen for the creation of logs.

Note: If multiple Data Manager DSNs refer to the same data store, they must all use exactly the same data store path name, even if some other path name identifies the same location. For example, you cannot use a symbolic link to refer to the data store in one DSN and the actual path name in another DSN. On Windows, you cannot use a mapped drive letter in the data store path name.

A Client DSN refers to a TimesTen data store indirectly by specifying a `hostname, DSN` pair, where the `hostname` represents a machine on which TimesTen Server Daemon is running and the `DSN` refers to a system DSN that is defined on that host. We refer to this host as the “server machine” and the `DSN` as a “Server DSN.”



On UNIX, all user DSNs (Client DSNs and/or Data Manager DSNs) created by a specific user are defined in the same user ODBC.INI file. Similarly, all system DSNs are defined in the same system ODBC.INI file.

The following table indicates the types of DSN supported by TimesTen, whether to create a user or system DSN and the location of the DSN.

DSN type	User or System DSN?	Location of DSN
Data Manager DSN	Can be a user or system DSN	Located on the machine where the data store resides.
Client DSN	Can be a user or system DSN	Located on any local or remote machine.
Server DSN	Must be a system DSN	Located on the machine where the data store resides.

The remainder of this chapter describes Data Manager DSNs and the connection attributes that can be defined for them. For more information about Client DSNs and Server DSNs, see [Chapter 2, “Working with the TimesTen Client and Server.”](#)

Connection attributes for Data Manager DSNs

There are four types of TimesTen Data Manager attributes:

- **Data store attributes** are associated with a data store when it is created and cannot be modified by subsequent connections.
- **First connection attributes** are set when a connection is made to an idle data store (a data store with no connections) and persists for that connection and all subsequent connections until the last connection to the data store is closed.
- **General connection attributes** are set by each connection and persist for the duration of the connection. Different concurrent connections may use different values.
- **Cache Connect attributes** allow you to enter the Oracle Service Identifier for the Oracle instance from which data will be loaded into TimesTen.

Note: See [Chapter 2, “Working with the TimesTen Client and Server”](#) for a description of the connection attributes that can be used with the TimesTen Client ODBC driver.

For a complete description of each attribute, see [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*.



On Windows, you specify data store attributes in the ODBC Data Source Administrator.

On UNIX, you specify data store attributes in the ODBC.INI file. Attributes that do not appear in the ODBC.INI file assume their default value.

Thread programming with TimesTen

TimesTen supports multi-threaded application access to data stores. When a connection is made to a data store, any thread may issue operations on the connection.

Typically, a thread issues operations on its own connection and therefore in a separate transaction from all other threads. In environments where threads are created and destroyed rapidly, better performance may be obtained by maintaining a pool of connections. Threads can allocate connections from this pool on demand to avoid the connect and disconnect overhead.

TimesTen allows multiple threads to issue requests on the same connection and therefore the same transaction. These requests are serialized by TimesTen, although the application may require additional serialization of its own.

TimesTen also allows a thread to issue requests against multiple connections, managing activities in several separate and concurrent transactions on the same or different data stores.

Creating a DSN on Windows



This section describes how to set up a TimesTen data store on Windows. Before you begin, read the [“TimesTen ODBC and JDBC drivers” on page 10](#) to find out what you’ll need to consider as you set up the data store.

For additional examples of setting up a data store, see [“DSN examples” on page 24](#).

This section includes the following topics:

- [Specify the ODBC driver](#)
- [Specify the DSN](#)
- [Specify the connection attributes](#)

Specify the ODBC driver

Specify the ODBC driver in the ODBC Data Source Administrator.

Note: JDBC users need to specify the ODBC driver to be used by the JDBC driver, as described in [“TimesTen JDBC driver” on page 12](#).

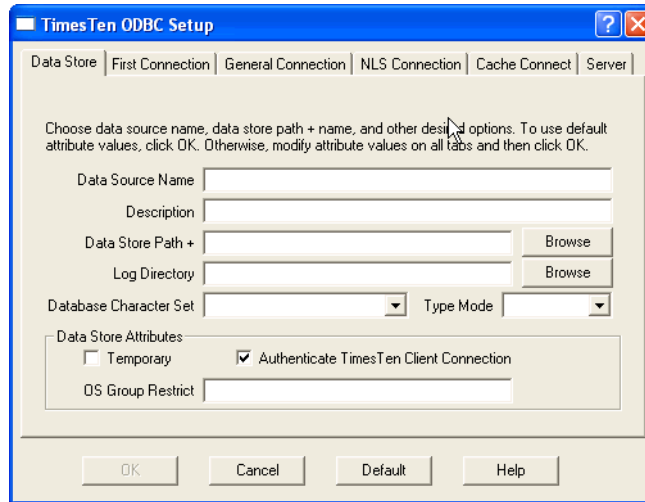
1. On the Windows Desktop, choose **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**. This opens the ODBC Data Source Administrator.
2. Choose **User DSN** if you want to create a user DSN or **System DSN** if you want to create a system DSN. You must have administrator privileges to create a System DSN. For a description of user and system DSNs, see [“User and system DSNs” on page 13](#).
3. Do one of the following:
 - Select an existing data source and click **Configure**.
 - Click **Add**, choose the appropriate TimesTen driver from the list. Click **Finish**. This displays the TimesTen ODBC Setup dialog.

For a list of TimesTen ODBC drivers, see [“TimesTen ODBC drivers” on page 10](#).

Specify the DSN

On the Data Store tab of the TimesTen ODBC Setup dialog, specify a data source name, a data store path name, and a database character set. The Data Store Path Name cannot reference a mapped drive. See [Figure 1.1](#).

Figure 1.1 Data Store



For an explanation of DSNs and data store path names, see [“Data source names” on page 13](#). For an explanation of database character sets, see [“Choosing a database character set” on page 76](#). The description field is optional.

Specify the connection attributes

Indicate the desired connection attributes under the **First Connection**, **General Connection**, and **NLS Connection** tabs of the TimesTen ODBC Setup dialog. See [Figure 1.2](#), [Figure 1.3](#), and [Figure 1.4](#). If you are using the Cache Connect for Oracle feature, specify the connection attributes shown in [Figure 1.5](#). If you are using a multithreaded client/server configuration, specify the connection attributes shown in [Figure 1.6](#).

For a description of the connection attributes, see [Chapter 1, “Data Store Attributes”](#) in *Oracle TimesTen In-Memory Database API Reference Guide*.

Figure 1.2 First Connection Attributes

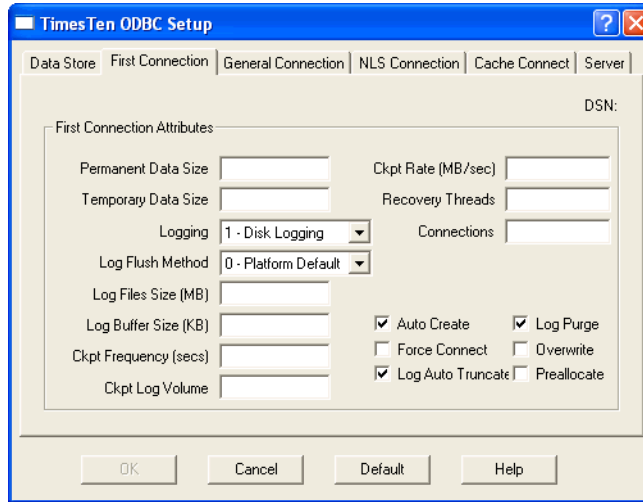


Figure 1.3 General Connection Attributes

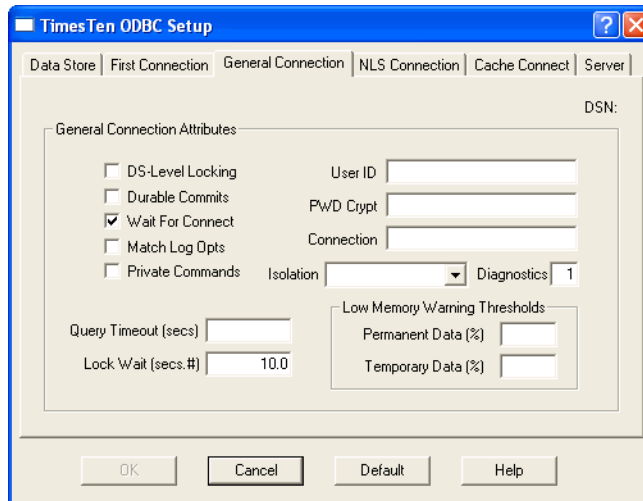


Figure 1.4 NLS Connection Attributes

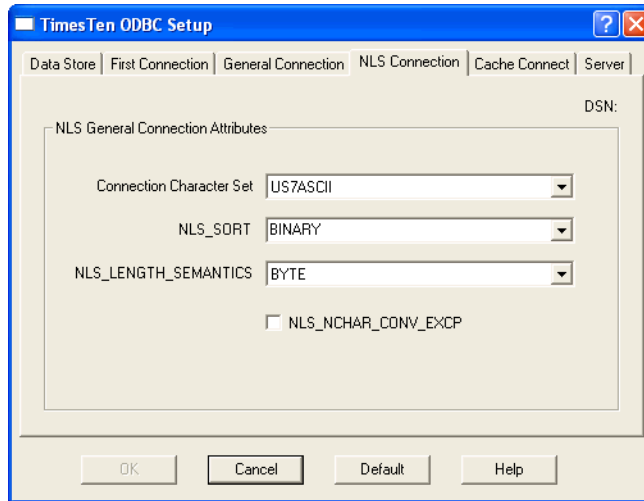


Figure 1.5 Cache Connect Attributes

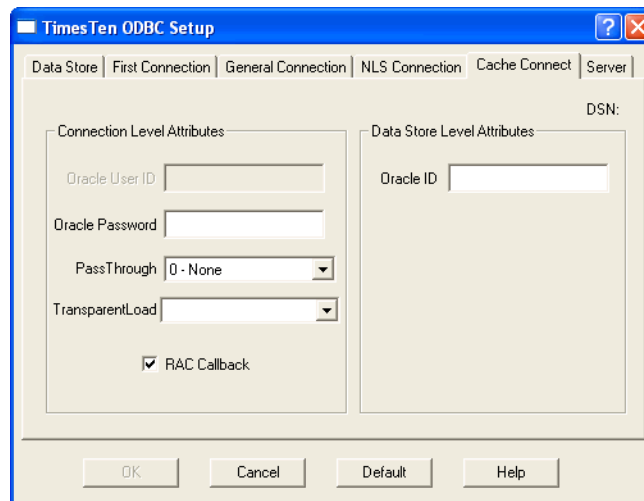
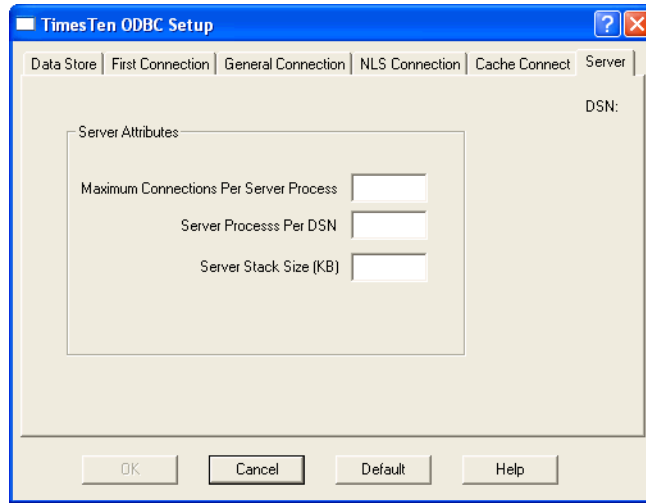


Figure 1.6 Server Attributes



Click **OK** when you are finished.

Creating a DSN on UNIX



This section describes how to set up a TimesTen data store on UNIX. Before you begin, be sure to read the [“TimesTen ODBC and JDBC drivers” on page 10](#) to find out what you’ll need to consider as you set up the data store.

For examples of defining a data store, see [“DSN examples” on page 24](#).

This section includes the following topics:

- [Create a user ODBC.INI file](#)
- [Specify the DSN](#)
- [Specify the ODBC driver](#)
- [Specify the data store path name](#)

Create a user ODBC.INI file

On UNIX, user DSNs are defined in the file `$HOME/.odbc.ini` or in a file named by the ODBCINI environment variable. This file is referred to as the “user ODBC.INI file.” System DSNs are defined in the file `/var/TimesTen/sys.odbc.ini` if TimesTen is installed by user `root`, or `install_dir/info/sys.odbc` if TimesTen is installed by a non-root user. This file is referred to as the “system ODBC.INI file.”

The syntax for a user ODBC.INI file and a system ODBC.INI file is the same.

The system ODBC.INI file is created when TimesTen is installed on the machine. Users must create their own user ODBC.INI file.

Specify the DSN

Specify the data source name in the ODBC.INI file. The DSN appears inside square brackets at the top of the DSN definition on a line by itself. For example:

```
[AdminDS]
```

Specify the ODBC driver

Note: JDBC users need to specify the ODBC driver to be used by the JDBC driver, as described in [“TimesTen JDBC driver” on page 12](#).

To set the TimesTen driver, specify the `DRIVER` parameter in the ODBC.INI file. For example:

```
[AdminDS]  
DRIVER=install_dir/lib/libtten.so
```

For a list of TimesTen ODBC drivers, see [“TimesTen ODBC drivers” on page 10](#).

Specify the data store path name

Specify the data store path name in the ODBC.INI file. For example:

```
DataStore=/users/robin/FixedDs
```

where FixedDs is the prefix for data store files. For more information, see [“Data source names” on page 13](#).

Choose a database character set

Specify a database character set in the ODBC.INI file. For example:

```
DatabaseCharacterSet=US7ASCII
```

For more information, see [“Choosing a database character set” on page 76](#).

Set data store attributes

Specify data store attributes in your ODBC.INI file. Attributes that do not appear in the ODBC.INI file assume their default value. See [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*. For examples, see [“DSN examples” on page 24](#).

Using environment variables in data store path names

You can use environment variables in the specification of the data store path name and log file path name. For example, you can specify `$HOME/AdminDS` for the location of the data store.

Environment variables can be expressed either as `$varname` or `$(varname)` (the parentheses are optional). A backslash character (`\`) in the data store path name quotes the next character.

Note: Environment variable expansion uses the environment of the process connecting to the data store. Different processes may have different values for the same environment variables and may therefore expand the data store path name differently. Environment variables can only be used in the user ODBC.INI file. They cannot be specified in the system ODBC.INI file.

DSN examples

This section provides additional examples of how to set up a data store:

- [Setting up a temporary data store](#)
- [Creating multiple DSNs to a single data store](#)
- [Connecting to a data store without using a DSN](#)

For each example, the Windows ODBC Data Source Administrator settings are followed by the corresponding ODBC.INI entries for UNIX.

Setting up a temporary data store

Example 1.1 This example illustrates how to set up a temporary data store.



On Windows, you can use the settings in the TimesTen ODBC Setup dialog to set up a temporary data store. See [Figure 1.7](#) and [Figure 1.8](#).

Figure 1.7 Data Store

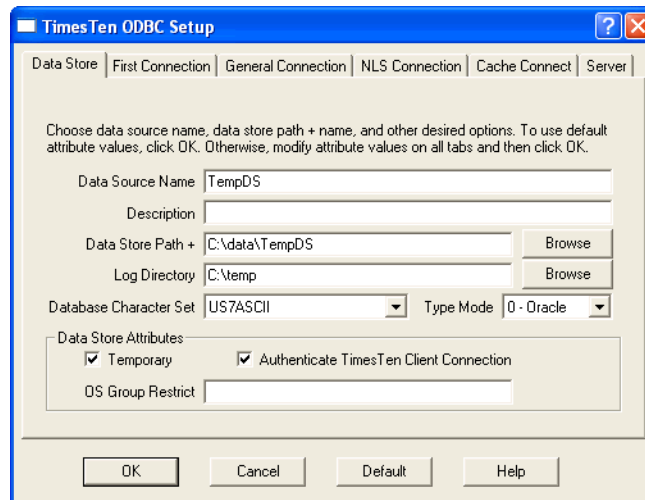
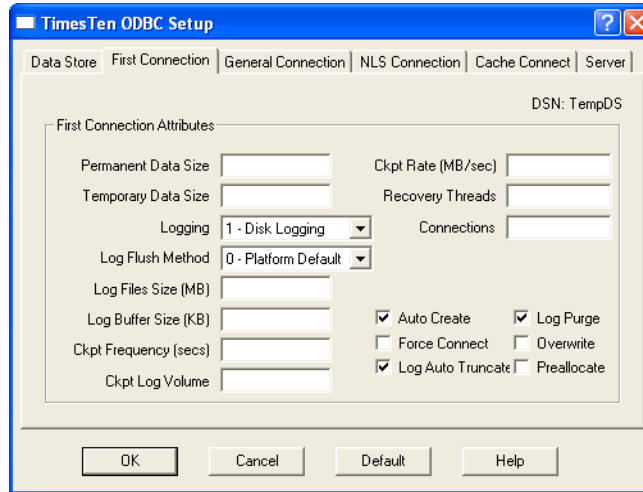


Figure 1.8 First Connection Attributes



To set up a temporary data store on UNIX, create the following entries in your ODBC.INI file. For a list of drivers for all UNIX platforms, see the table on page 11.

The text in square brackets is the data source name.

```
[TempDs]
Driver=install_dir/lib/libtten.so
DataStore=/users/robin/TempDs
#this is a temporary data store
Temporary=1
#create data store if it is not found
AutoCreate=1
#log data store updates to disk
Logging=1
LogPurge=1
DatabaseCharacterSet=US7ASCII
```

Note: A temporary data store cannot be backed up.

Creating multiple DSNs to a single data store

You can create two or more DSNs that refer to the same data store but have different connection attributes.

Example 1.2 This example creates two DSNs, AdminDSN and GlobalDSN. The DSNs are identical except for their connection character sets. Applications that use the US7ASCII character set can connect to the TTDS data store by using AdminDSN.

Applications that use multibyte characters can connect to the TTDS data store by using GlobalDSN.



For Windows, use the ODBC Data Source Administrator to define AdminDSN as shown in Figure 1.9. AdminDSN is created with the AL32UTF8 database character set. Figure 1.10 shows that US7ASCII is the connection character set for AdminDSN.

Figure 1.9 Creating AdminDSN using TTDS data store

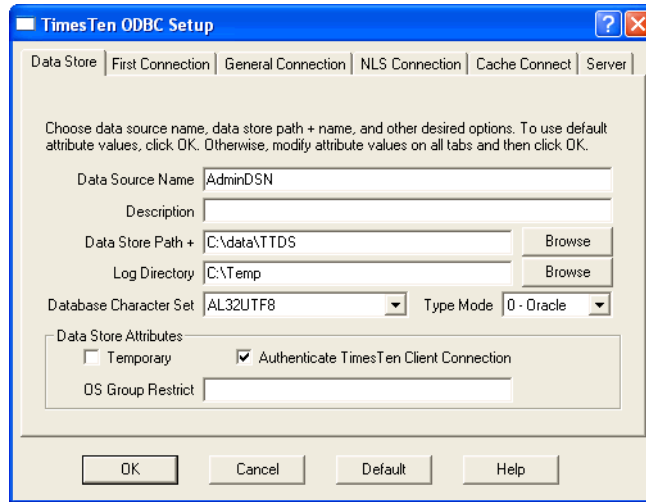
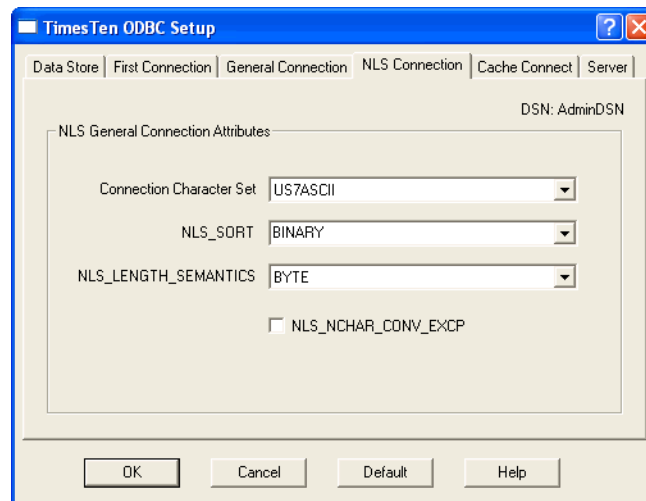


Figure 1.10 Setting the connection character set for AdminDSN



GlobalDSN is also created with the AL32UTF8 database character set, as shown in Figure 1.11. Figure 1.12 shows that the connection character set for GlobalDSN is AL32UTF8.

Figure 1.11 Creating GlobalDSN using TTDS data store

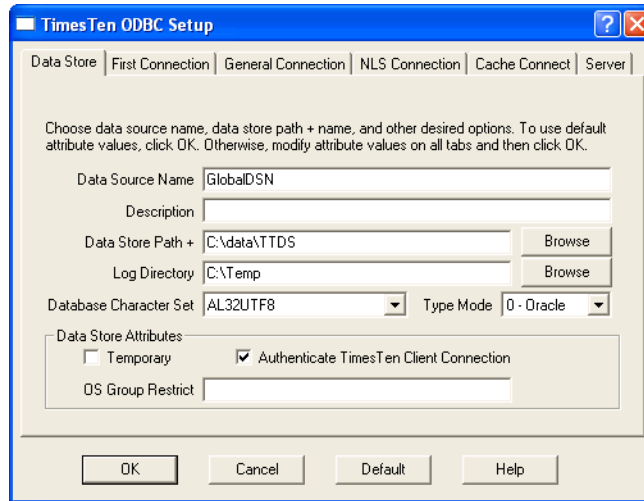
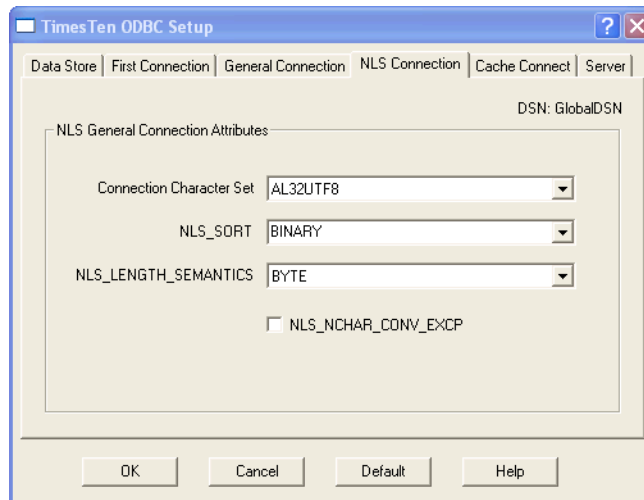


Figure 1.12 Setting the connection character set for GlobalDSN



Example 1.3 shows how to specify the DSNs on UNIX. It uses the TimesTen Data Manager ODBC driver for Solaris.

Example 1.3 The text in square brackets is the data source name.

```
[AdminDSN]
Driver=install_dir/lib/libtten.so
Datastore=/data/TTDS
DatabaseCharacterSet=AL32UTF8
ConnectionCharacterSet=US7ASCII
```

```
[GlobalDSN]
Driver=install_dir/lib/libtten.so
DataStore=/data/TTDS
DatabaseCharacterSet=AL32UTF8
ConnectionCharacterSet=AL32UTF8
```

Connecting to a data store without using a DSN

Using the **ttIsql** utility, you can connect to a data store without a predefined Data Source Name by specifying:

- The name or path name of driver using the Driver attribute, and
- The data store path and filename prefix using DataStore attribute

On Microsoft Windows systems, the value of the Driver attribute should be the name of the TimesTen ODBC Driver. For example, TimesTen Data Manager 7.0.

On UNIX systems, the value of the Driver attribute should be the pathname of the TimesTen ODBC Driver shared library file. The file resides in the *install_dir/lib* directory.

Example 1.4 C:\ ttIsql

```
ttIsql <c> 1996-2006, Oracle. All rights reserved.
Type ? or "help" for help, type "exit" to quit ttIsql.
All commands must end with a semicolon character.

Command> connect "Driver=TimesTen Data Managers7.0;
                DataStore=C:\sales\admin";
```

Specifying the size of a data store

This section includes the following topics:

- [Temporary and permanent memory](#)
- [Changing data store size](#)
- [Receiving out-of-memory warnings](#)

Temporary and permanent memory

TimesTen manages data store space using two separate memory partitions within a single contiguous memory space. One partition contains permanent data and the other contains temporary data.

- Permanent data includes the tables and indexes that make up a TimesTen data store. When a data store is loaded into memory, the contents of the permanent data partition are read from files stored on disk. The permanent data partition is written to disk during checkpoint operations.
- Temporary data includes locks, cursors, compiled commands, and other structures needed for command execution and query evaluation. The temporary data partition is created when a data store is loaded into memory and is destroyed when it is unloaded.

The connection attributes that control the size of the data store when it is in memory are `PermSize` and `TempSize`. The `PermSize` attribute specifies the size of the permanent data partition and the `TempSize` attribute specifies the size of the temporary data partition.

See [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide* for further description of these attributes.

Changing data store size

The sizes of the permanent and temporary data partitions are set when a data store is loaded into memory and cannot be changed while the data store is in memory. To change the size of either partition, you must unload the data store from memory and then reconnect using different values for the `PermSize` or `TempSize` attributes. You can use the `ttStatus` utility to find processes connected to the data store and stop them. Once you have made the change in data store size, reload it into memory.

If the data store is configured for replication, you must also stop the replication agent. You must reconfigure the data store sizes for all replicas of the data store. Once you have made the change in data store size, read it into memory and restart the replication agent and Cache Connect if it was connected.

- The permanent data partition can be increased in size, but it cannot be decreased.

- The temporary data partition can be either increased or decreased in size for data stores that do not participate in replication.

To unload a data store from memory, you must close all active connections to the data store and the RAM policy of the data store must be set to manual or inUse. See “[Specifying a RAM policy](#)” on page 31 for more information about RAM policy settings.

You must also make sure that you have a shared memory segment that is large enough to hold the data store. In general, the minimum size of this shared memory segment should be:

$$\text{PermSize} + \text{TempSize} + \text{LogBuffSize} + 7\text{MB overhead}$$

For more details, see “[Installation prerequisites](#)” in *Oracle TimesTen In-Memory Database Installation Guide* and the descriptions of the attributes **TempSize** and **PermSize** in *Oracle TimesTen In-Memory Database API Reference Guide*.

Receiving out-of-memory warnings

TimesTen provides two General Connection attributes that determine when a low memory warning should be issued: **PermWarnThreshold** and **SqlQueryTimeout**. Both attributes take a percentage value.

To receive out-of memory warnings, applications must call the built-in procedure **ttWarnOnLowMemory**.

These attributes also set the threshold for SNMP warning. See the chapter, “[Diagnostics through SNMP Traps](#)” in the *Oracle TimesTen In-Memory Database API Reference Guide*.

Specifying a RAM policy

TimesTen allows you to specify a RAM policy that determines when data stores are loaded and unloaded from main memory. To set the RAM policy, use the [ttAdmin](#) utility.

For each data store you can have a different RAM policy. The policy options are:

- *In Use*. The data store is loaded into memory when the first connection to the data store is opened, and it remains in memory as long as it has at least one active connection. When the last connection to the data store is closed, the data store is unloaded from memory. This is the default policy.
- *InUse with RamGrace*. The data store is loaded into memory when the first connection to the data store is opened, and it remains in memory as long as it has at least one active connection. When the last connection to the data store is closed, the data store remains in memory for a “grace period.” The data store is unloaded from memory only if no processes have connected to the data store for the duration of the grace period. The grace period can be set or reset at any time. It is only in effect and stays in effect until the next time the grace period is changed.
- *Always*. The data store stays in memory at all times. If the machine on which the data store resides is rebooted, the data store reloads into memory when the TimesTen daemon is started, generally at boot time.
- *Manual*. The data store is manually loaded and unloaded by system administrators using the [ttAdmin](#) utility.

Copying, migrating, backing up and restoring a data store

The TimesTen utilities for copying, backing up, restoring and migrating a data store allow you to:

- Migrate a data store between releases of TimesTen
- Migrate a data store between different hardware platforms
- Add rows of data to a table
- Take a snapshot of a data store and then restore it
- Rename the owner of tables in a data store

To migrate a data store between releases of TimesTen, use the **ttMigrate** utility. This utility saves tables and indexes from a TimesTen data store into a binary file. The tables and indexes can then be restored into another TimesTen data store. This allows you to migrate data between TimesTen releases.

To migrate a data store between hardware platforms, use the **ttBulkCp** utility. This utility saves the rows of a table to an ASCII file. It allows you to copy a single table between data stores, including between data stores from different releases of TimesTen or between data stores on different hardware platforms.

To add rows of data to an existing table, use the **ttBulkCp** utility. You can save data to an ASCII file and use the **ttBulkCp** utility to load the data rows into a table in a TimesTen data store. The rows you are adding must contain the same number of columns as the table, and the data in each column must be of the type defined for that column. Because the **ttBulkCp** utility works on data stored in ASCII files, you can also use this utility to import data from other applications, provided the number of columns and data types are compatible with those in the table in the TimesTen data store and that the file found is compatible with **ttBulkCp**.

To take a snapshot of a data store and later restore that data store in the exact same state, use the **ttBackup** and **ttRestore** utilities or the **ttBackup** and **ttRestore** utilities C functions.

To rename the owner of tables in a data store, use the **ttMigrate** utility. When restoring tables, you can use the `-rename` option to rename the owner of tables.

Backing up and restoring a data store

TimesTen's backup and restore facility allows you to create backups of TimesTen data stores and restore the data store at a later time. The primary use for the backup and restore facility is to allow the restoration of a recent state of a data store that has been lost.

Every data store backup contains the information needed to restore the data store as it existed at a the *backup point*; the time the backup began. Restoration of a data store from a given backup restores the modifications of all transactions that committed before the backup point.

In this release, TimesTen supports both *full* and *incremental* backups. An incremental backup moves the backup point of an existing backup forward in time by augmenting the backup with all of the log records created since its backup point.

TimesTen writes a data store backup to a location specified by a *backup path*, which consists of a *directory name* and an optional *basename*. You must specify the backup directory and basename when the backup is created. The basename defaults to the basename of the data store itself if you do not specify a basename.

Note: You must not manually change the contents of the backup directory. The addition, removal, or modification of any file in the backup directory, except for modifications made by ttBackup and ttRestore themselves, may compromise the integrity of the backup and restoration of the data store from the backup may not be possible.

TimesTen also allows *stream* backups. A stream backup writes the data store backup file to `stdout`.

A set of files containing backup information for a given data store, residing at a given backup path is referred to as a *backup instance*. A given backup instance must be explicitly enabled for incremental backups.

An incremental backup can only augment an existing *incremental-enabled* backup of the same data store. Restoring a data store from a backup causes all existing incremental-enabled backups of this data store to become incremental backups to become incremental-disabled.

TimesTen supports the creation of up to eight incremental-enabled backup instances for each data store. If you attempt to start an incremental backup in a ninth backup path, TimesTen returns an error. Incremental backups are supported only for permanent disk-logging data stores.

Information about incremental backups is not retained across data store recovery or restoration. Existing backup instances continue to be usable for restoring the data store, but incremental backups to those instances will have to be reenabled.

A backup operation is atomic: If it completes successfully, it will produce a backup that can be used to restore a data store to the state of its backup point. If it fails for any reason, it leaves the files of the existing backup (if any) intact and its backup point unchanged.

Note: For full backups, you must have enough disk space available to hold both the existing backup and the new backup, until the new backup succeeds.

The files of the existing backup may be modified by a failed full or incremental backup, but not in a way that compromises the ability to restore from them.

An incremental backup typically completes much faster than a full backup, as it has less data to copy. The performance gain of incremental backups over full backups comes at the cost of increased disk usage and longer restoration times. Use incremental backups in concert with full backups in order to achieve a balance between backup time, disk usage, and restoration time.

The backup types supported by TimesTen are:

Backup Type	File or Stream	Full or Incremental	Incremental -enabled	Comments
fileFull	File	Full	No	Default
fileFullEnable	File	Full	Yes	
fileIncremental	File	Incremental.	Yes	Fails if incremental backup not possible.
fileIncrOrFull	File	Either	Yes	Performs fileIncremental if possible; fileFullEnable otherwise.
streamFull	Stream	Full	No	
incrementalStop	None	None	No	Takes no backup; just disables existing incremental-enabled backup.

For details on using TimesTen's backup and restore facility, see these references in the [Oracle TimesTen In-Memory Database API Reference Guide](#).

Working with the ODBC.INI file

This section includes the following topics:

- [The user ODBC.INI file](#)
- [The system ODBC.INI file](#)
- [Searching for a DSN](#)
- [ODBC Data Sources](#)
- [Data Source Specification](#)
- [odbc.ini file example](#)

The user ODBC.INI file

On UNIX, user DSNs are defined in the file `$HOME/.odbc.ini` or in a file named by the ODBCINI environment variable. This file is referred to as the “user ODBC.INI file.” Although a user DSN is private to the user who created it, it is only the DSN (the character-string name and its attributes) that is private. The underlying data store can be referenced by other user DSNs or by system DSNs. TimesTen supports data sources for the TimesTen Data Manager and data sources for the TimesTen Client in the `.odbc.ini` file.

For information on how to create a copy of the `.odbc.ini` file in your home directory and how to override the name and location of the `.odbc.ini` file, see [“Data source names” on page 13](#).

The system ODBC.INI file

On UNIX, system DSNs are defined in the `/var/TimesTen/sys.odbc.ini` file. This file is referred to as the “system ODBC.INI file.” A system DSN can be used by any user on the machine on which the system DSN is defined.

Searching for a DSN

See [“Searching for a DSN” on page 52](#) for the rules of precedence that TimesTen follows when searching for a DSN.

ODBC Data Sources

Each entry in the optional ODBC Data Sources section lists a data source and a description of the driver it uses. The data source section has the following format:

```
[ODBC Data Sources]
data-source-name=driver-description
```

The *data-source-name* is required. It identifies the data source to which the driver connects. You choose this name.

The *driver-description* is required. It describes the driver that connects to the data source.

Data Source Specification

Each data source listed in the ODBC Data Sources section has its own data source specification section. The data store specification for TimesTen Data Manager data stores has the format shown in [Table 1.1](#).

Note: These examples reference the 32-bit sample DSNs. This is indicated by the extension `_32`. On 64-bit platforms, the sample DSNs are appended with `_64`.

Table 1.1 Data Source Specification Format

Component	Description
[<i>data-source-name</i>]	The <i>data-source-name</i> is required. It is the name of the data source, as specified in the ODBC Data Sources section of your <code>.odbc.ini</code> file.
Driver= <i>driver-path-name</i>	The TimesTen Data Manager driver that is linked with the data source.
DataStore= <i>data-store-path-name</i>	The path name of the data store to access. The path name is required.
Optional attributes	See " Data Store Attributes " for information about attributes.

For example, a data source called **RunData_tt70_32** may have the following data source specification entry:

```
[RunData_tt70_32]
Driver=install_dir/lib/libtten.sl
DataStore=/users/robin/SalesDs
#create data store if it is not found
AutoCreate=1
#do not wait if cannot connect to data store
WaitForConnect=0
#remove old log files at connect and checkpoint
LogPurge=1
```

The data store specification for TimesTen Client configurations has the format shown in [Table 1.2](#).

Table 1.2 Data Store Specification for TimesTen Client Configurations

Component	Description
<code>[data-source-name]</code>	The <i>data-source-name</i> is required. It is the name of the data source, as specified in the ODBC Data Sources section of your <code>.odbc.ini</code> file.
<code>TTC_Server=server-name</code>	The <i>server-name</i> is required. It is the DNS name, host name, IP address or shorthand name for the TimesTen Server.
<code>TTC_Server_DSN=server-DSN</code>	The <i>server-DSN</i> is required. It is the name of the data source to access on the TimesTen Server.
<code>TTC_Timeout=value</code>	Client connection timeout value in seconds.

Note: Most TimesTen Data Manager attributes are ignored for TimesTen Client data stores.

For example, the data source called **RunDataCS_tt70_32** that connects to the data source, named **RunData_tt70_32**, on the *ttserver* TimesTen Server may have the data source specification entry:

```
[RunDataCS_tt70_32]
TTC_Server=ttserver
TTC_Server_DSN=RunData_tt70_32
TTC_Timeout=30
```

For example, the data source called **ShmDataCS_tt70_32** that uses a shared memory segment to connect to the data source named **RunData_tt70_32** on the *ShmHost70* TimesTen Server may have the data source specification entry:

```
[ShmRunDataCS_tt70_32]
TTC_Server=ShmHost70
TTC_Server_DSN=RunData_tt70_32
TTC_Timeout=30
```

odbc.ini file example

The following example shows a UNIX `.odbc.ini` file:

```
[ODBC Data Sources]
RunData_tt70_32=TimesTen 7.0 Driver
RunDataCS_tt70_32=TimesTen Client 7.0
```

```
[RunDataCS_tt70_32]
TTC_Server=tt_server_logical
TTC_Server_DSN=RunData
TTC_Timeout=30

[RunData_tt70_32]
Driver=install_dir/lib/libtten.sl
DataStore=/users/robin/RunData
PermSize=8
Logging=1
```


Working with the TimesTen Client and Server

To access TimesTen data stores on remote machines, applications can use the TimesTen ODBC or JDBC Client driver and TimesTen Server Daemon. Using the TimesTen Client, applications written to use the TimesTen Data Manager can connect transparently to TimesTen data stores on any remote or local machine that has the TimesTen Server Daemon and Data Manager installed.

You can also install the TimesTen Client and TimesTen Server on the same machine and use them to access TimesTen data stores on the local machine. On Solaris, HP-UX and other platforms that support both 32-bit and 64-bit applications, this is useful if you have a 32-bit application that needs to access a 64-bit data store on the same machine. You can create a client/server connection between any combination of 32/64 bits RedHat Linux, Solaris, and HP-UX and Windows platforms.

The TimesTen Server is a process that runs on a server machine. The TimesTen ODBC or JDBC Client is a thin driver that communicates with the TimesTen Server.



On Windows, you can either link the Client applications with the driver manager or directly with the TimesTen Client driver.



On UNIX, you must link Client applications directly with the TimesTen Client driver.

For UNIX, there are Client/Server versions of some TimesTen utilities: ttIsqlCS, ttBulkCpCS, ttMigrateCS and ttSchemaCS.

For details on compiling TimesTen applications, see the [Oracle TimesTen In-Memory Database Java Developer's and Reference Guide](#) or the [Oracle TimesTen In-Memory Database C Developer's and Reference Guide](#).

The main topics in this chapter are:

- [Client/Server Communication](#)
- [Configuring TimesTen Client and Server](#)
- [Running the TimesTen Server Daemon](#)
- [TimesTen Server connection attributes](#)
- [Defining Server DSNs](#)

- [TimesTen Client connection attributes](#)
- [Creating and configuring Client DSNs on Windows](#)
- [Creating and configuring Client DSNs on UNIX](#)
- [Accessing a remote data store on UNIX](#)
- [Working with the TTCONNECT.INI file](#)

Client/Server Communication

Each TimesTen Client connection requires one Server process. By default, a Server process is spawned at the time a Client requests a connection. By setting the `-serverPool` option in the `ttendaemon.options` file on the Server machine, you can pre-spawn a reserve pool of Server processes. See [“Prespawning TimesTen Server processes” on page 68,](#)” for details.

When using TimesTen Client/Server there are three ways the TimesTen Client can communicate with the TimesTen Server.

TCP/IP Communication

By default, the TimesTen Client communicates with the TimesTen Server daemon using TCP/IP sockets. This is the only form of communication available when the TimesTen Client and Server are installed on different machines.

Shared memory communication

If both the TimesTen Client and Server are installed on the same machine, applications using the TimesTen Client ODBC driver may improve performance by using a shared memory segment for inter-process communication (IPC). Using a shared memory segment allows for the best performance, but consumes more memory. To use a shared memory segment as communication, you must set the server options in the `ttendaemon.options` file. See [“Using shared memory for Client/Server IPC” on page 69.](#) You must also define the Network Address of the logical server as `ttShmHost`. See [“Creating and configuring Client DSNs on Windows” on page 45](#) or [“Creating and configuring Client DSNs on UNIX” on page 51.](#)

Note: TimesTen supports a maximum of 16 different instances of the shared memory IPC-enabled server. If an application tries to connect to more than 16 different shared memory segments it receives the ODBC error “Cannot connect to more than 16 SHMIPC-enabled TimesTen Servers.”

UNIX domain socket communication

On certain UNIX platforms, if both the TimesTen Client and Server are installed on the same machine, you can use UNIX domain sockets for communication.

Using a shared memory segment allows for the best performance, but greater memory usage. Using UNIX domain sockets allows for improved performance over TCP/IP, but with less memory consumption than a shared memory segment connection. See the section in this chapter that describes how to create a Client DSN for your platform. To use domain sockets, you must define the Network Address of the logical server as `ttLocalHost`. See [“Creating and configuring Client DSNs on UNIX” on page 51](#).

Configuring TimesTen Client and Server

Before configuring the TimesTen Client and Server, read the ["TimesTen ODBC and JDBC drivers"](#) and ["Data source names"](#) sections of [Chapter 1, “Creating TimesTen Data Stores.”](#) To connect a TimesTen application to a TimesTen data store using TimesTen Client and Server:

Configuring Client/Server of the same TimesTen release

To configure TimesTen when the Client and Server are of the same TimesTen release:

1. Install the TimesTen Server on the machine on which the TimesTen data store resides. This machine is called the “server machine.” For information on how to install the TimesTen Server, see the [Oracle TimesTen In-Memory Database Installation Guide](#). During installation, choose to have the TimesTen Server automatically started.
2. Install the TimesTen Client on the machine where the Client application resides. This machine is called the “client machine.” For information on how to install the TimesTen Client, see the [Oracle TimesTen In-Memory Database Installation Guide](#).
3. For JDBC, install the Java Developer’s Kit (JDK) on the client machine, where the Java application(s) will be running and set up the environment variables (CLASSPATH and LIBRARY PATH) on the client machine. See [Chapter 1, “Setting the Java environment variables in the Oracle TimesTen In-Memory Database Java Developer’s and Reference Guide](#) for details.
4. On the server machine, create and configure a Server DSN corresponding to the TimesTen data store. See [“Defining Server DSNs” on page 43](#).
5. On the client machine, create and configure a Client DSN corresponding to the Server DSN. See [“Creating and configuring Client DSNs on UNIX” on page 51](#) and [“Creating and configuring Client DSNs on Windows” on page 45](#).
6. Set TimesTen Client connection attributes. See [Chapter 1, “Data Store Attributes”](#) in the [Oracle TimesTen In-Memory Database API Reference Guide](#).
7. For C client applications, link as described in ["Linking options"](#) in the [Oracle TimesTen In-Memory Database C Developer’s and Reference Guide](#).

Configuring cross-release Client/Server

A TimesTen 6.0 client can connect to a 6.0 TimesTen Server of any patch level. If the `-insecure-backwards-compat` option is set in the `ttendaemon.options` file for a TimesTen 7.0 or newer instance, a TimesTen 6.0 client can connect to a TimesTen 7.0 or newer server, under certain conditions and a TimesTen 7.0 or later client can connect to a TimesTen 6.0 or newer server, under certain these conditions:

- You must install the TimesTen Data Manager on the machine where the Server is installed.
- There must be a TimesTen daemon running on the Server machine whose version matches exactly the version of the TimesTen Server.
- There must be a TimesTen daemon running on the Server machine whose version matches exactly the version of the TimesTen Data Manager.
- The driver of the Data Manager should be specified in the Server DSN.
- The driver of the Data Manager must be of the same bit-level (32 or 64) as the bit-level of the Server.
- The version of the driver of the Data Manager must match exactly the version of the data store.

Running the TimesTen Server Daemon

The TimesTen Server daemon is a subdaemon of the TimesTen daemon. If the TimesTen Server is installed on your machine, the TimesTen Server daemon starts automatically every time the TimesTen daemon is started and stops automatically every time the TimesTen daemon is stopped on your machine.

The TimesTen Server daemon handles request from applications linked with the TimesTen Client ODBC driver.

By default, the 32-bit version of TimesTen Server Daemon listens on TCP port number 17002 and the 64-bit version listens on 17003. System administrators can change the port number during installation to avoid conflicts or for security reasons. The port range is from 1 - 65535. To connect to the TimesTen Server Daemon, Client DSNs are required to specify the port number as part of the logical server name definition or in the connection string.

On Windows, the TimesTen Server daemon is run as user SYSTEM. On UNIX, the TimesTen Server daemon is run as the instance administrator (user `root`, unless the instance is installed as non-root).

For instructions on modifying TimesTen Server Daemon options, see "[Modifying the TimesTen Server daemon options](#)" in Chapter 3 of the *Oracle TimesTen In-Memory Database Operations Guide*.

Server informational messages

The TimesTen Server records “connect,” “disconnect” and various warning, error and informational entries in log files.



On Windows, these application messages can be accessed with the Event Viewer.

On UNIX, the TimesTen Server logs messages either to files or to the `syslog` facility.

To specify using `syslog`, you set the `-userlog` or `-supportlog` option in the `ttendaemon.options` file to `syslog`. You can define the facility used by `syslog` by setting the `-facility` option in the `ttendaemon.options` file. Otherwise, the values specified to the `-userlog` or `-supportlog` options are the names of the files to be used for user messages or support messages, respectively. See [“Modifying informational messages” on page 64](#).

The `syslog` facility allows messages to be routed in a variety of ways, including recording them to one or more files. The disposition of messages is under the control of the configuration file `/etc/syslog.conf`. See the operating system documentation for `syslog.conf` or `syslogd` for information on how to configure this file. Message files can grow to be quite large. Prune them periodically to conserve disk space.

TimesTen Server connection attributes

By default, TimesTen creates only one connection to a Server per child process.

Because a TimesTen Server DSN corresponds to data stores that are accessed by a TimesTen Server daemon, a TimesTen Server DSN can be configured using regular connection attributes. In addition, you can specify connection attributes that allow you to specify multiple client connections to a single Server.

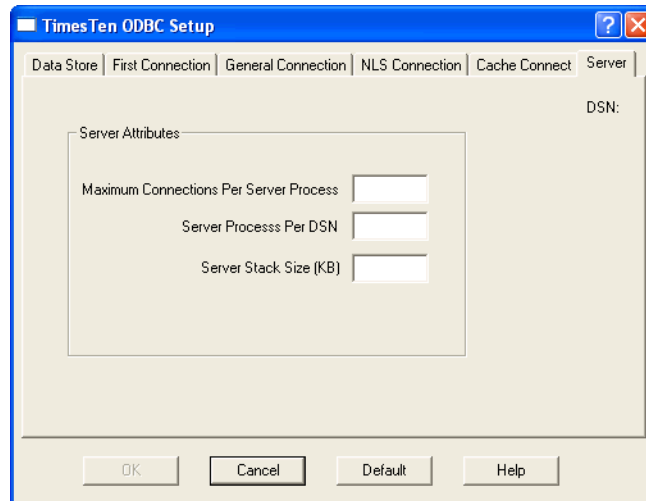
You can configure these attributes as part of the Server DSN definition or you can configure these attributes in the TimesTen daemon options file (`ttendaemon.options`). For a description of the TimesTen daemon options see [“Specifying multiple connections to the TimesTen Server” on page 68](#).

The Server Connection attributes are **MaxConnsPerServer**, **ServersPerDSN**, and **ServerStackSize**. For a complete description of the TimesTen Server connection attributes, see [Chapter 1, “Data Store Attributes” in the *Oracle TimesTen In-Memory Database API Reference Guide*](#).

Defining Server DSNs

Server DSNs correspond to data stores that are accessed by a TimesTen Server Daemon. You can add or configure a Server DSN while the TimesTen Server is running.

A Server DSN is a TimesTen data manager system DSN. For a description of DSNs and instructions on creating them, see [“Creating a DSN on Windows” on page 17](#) or [“Creating a DSN on UNIX” on page 22](#). If you anticipate having more than one connection to the Server DSN, specify appropriate values for the Server attributes as needed. On Windows, these are specified on the **Server** tab.



TimesTen Client connection attributes

This section discusses the connection attributes used by the TimesTen Client driver. You can configure attributes as part of the Client DSN definition or you can configure connection attributes at runtime in the *connection* string that is passed to the ODBC **SQLDriverConnect** function or the *URL* string that is passed to the JDBC **DriverManager.getConnection** method.

The **Authenticate** and **DataStore** Data Manager attributes are not allowed in the TimesTen Client connection string.

Note: The TimesTen Client allows TimesTen Data Manager attributes to be passed in as part of the connection or URL string. However, the Client ignores any Data Manager attributes specified in the ODBC.INI file as part of the TimesTen Client DSN definition.

For a complete description of the TimesTen Client connection attributes, see [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*.

Creating and configuring Client DSNs on Windows



On Windows, use the ODBC Data Source Administrator to configure logical server names and to define Client DSNs.

This section includes the following topics:

- [Creating and configuring a logical server name](#)
- [Creating a Client DSN on Windows](#)
- [Setting the timeout interval and authentication](#)
- [Deleting a server name](#)
- [Accessing a remote data store on Windows](#)
- [Testing connections](#)

Creating and configuring a logical server name

To create and configure a logical server name:

1. On the Windows Desktop, choose **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**.

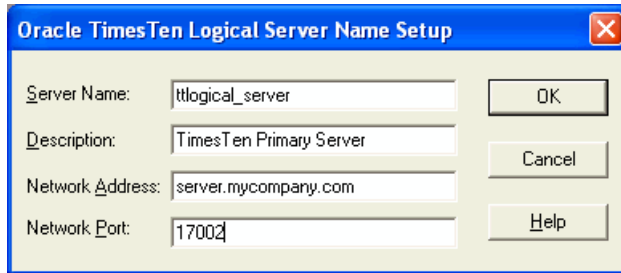
This opens the ODBC Data Source Administrator.

2. Click **User DSN** or **System DSN**.
3. Select a TimesTen Client DSN and click **Configure**. If no Client DSN exists, click **Add**, select **TimesTen Client 7.0** and click **Finish**. This opens the TimesTen Client DSN Setup dialog.
4. Click **Servers**. This opens the TimesTen Logical Server List dialog.
5. Click **Add**. This opens the TimesTen Logical Server Name Setup dialog.
6. In the **Server Name** field, enter a logical server name.
7. In the **Description** field, enter an optional description for the server.
8. In the **Network Address** field, enter the host name or IP address of the server machine. The Network Address must be one of:

Type of Connection	Network Address
Local Client/Server connection that uses shared memory for inter-process communication	tt.ShmHost
Remote Client/Server connection	The name of the machine where the TimesTen Server is running. For example, <code>server.mycompany.com</code>

- In the **Network Port** field, TimesTen displays the port number on which the TimesTen Logical Server Listens by default. If the TimesTen Server is listening on a different port, enter that port number in the **Network Port** field.

For example:



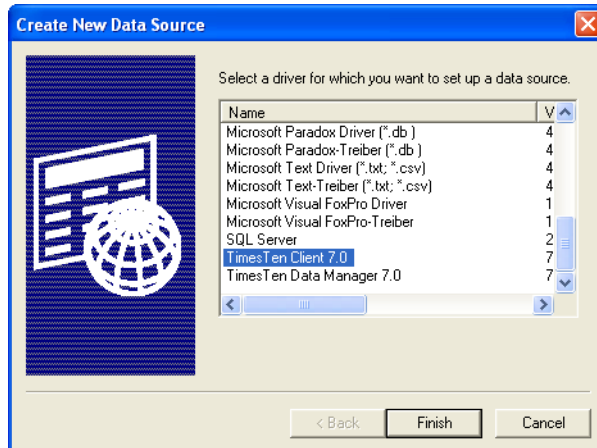
- Click **OK**, then click **Close** in the TimesTen Logical Server List dialog to finish creating the logical server name.

Creating a Client DSN on Windows

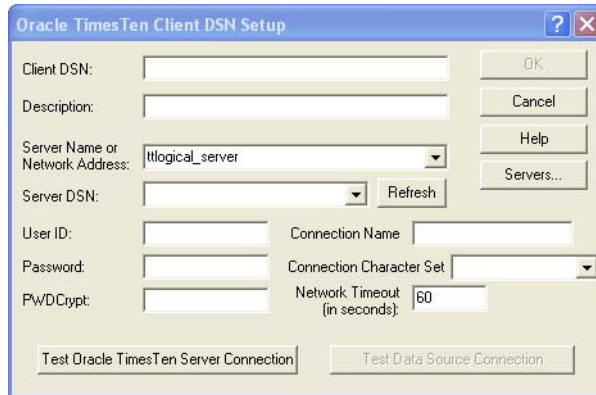


To define a TimesTen Client DSN:

- On the Windows Desktop, choose **Start > Settings > Control Panel > Administrative Tools > Data Sources (ODBC)**. This opens the ODBC Data Source Administrator.
- Choose either **User DSN** or **System DSN**. For a description of User DSNs and System DSNs see [“Data source names” on page 13](#).
- Click **Add**. This opens the Create New Data Source dialog.



4. Choose **TimesTen Client 7.0**. Click **Finish**. This opens the TimesTen Client DSN Setup dialog.



5. In the **Client DSN** field, enter a name for the Client DSN.

The name must be unique to the current list of defined DSNs on the machine where the Client application resides and can contain up to 32 characters. To avoid potential conflicts, you may want to use a consistent naming scheme that combines the logical server name with the name of the Server DSN. For example, a corporation might have Client DSNs named `Boston_Accounts` and `Chicago_Accounts` where `Boston` and `Chicago` are logical server names and `Accounts` is a Server DSN.
6. In the **Description** field, enter an optional description for the Client DSN.
7. In the **Server Name or Network Address** field, specify the logical server or network address of the server machine.
 - The name can be a host name, IP address or logical server name. The logical server names defined on the client machine can be found in the drop-down list. To define logical server names, click **Servers**.
 - If you do not specify a logical server name in this field, the TimesTen Client assumes that the TimesTen Server Daemon is running on the default TCP/IP port number. Therefore, if your Server is running on a port other than the default port and you do not specify a logical server name in this field, you must specify the port number in the ODBC connection string, using the `TCP_Port` attribute.

For more information on defining logical server names, see [“Creating and configuring a logical server name” on page 45](#).
8. In the **Server DSN** field, enter the Server DSN corresponding to the data store that the Client application will access.

- If you do not know the name of the Server DSN, click **Refresh** to obtain a list of Server DSNs that are defined on the machine specified in the **Server Name** or **Network Address** field. Select the Server DSN from the drop-down list.
 - You must have a network connection to the machine where the TimesTen Server is running.
9. In the **Connection Character Set** field, choose a character set that matches your terminal settings or your data source. The default connection character set is US7ASCII. For more information, see "[ConnectionCharacterSet](#)" in *Oracle TimesTen In-Memory Database API Reference Guide*.

Setting the timeout interval and authentication

For a description of the Timeout, UID and PWD attributes, see [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*.

To set the timeout interval and authentication:

1. In the **User ID** field of the Oracle TimesTen Client DSN Setup dialog box, you can enter a user name that is defined on the server machine. If the server DSN that corresponds to this client DSN is defined with `Authenticate=1`, you must provide a user name either in this field or in the connection string of every application that uses this client DSN. If the server DSN is defined with `Authenticate=0`, then TimesTen ignores the value entered in this field, unless Access Control is enabled. In that case the User ID is required.
2. In the **Password** field, you can enter the password that corresponds to the user ID. Alternatively, you can enter an encrypted password in the **PwdCrypt** field. If the server DSN that corresponds to this client DSN is defined with `Authenticate=1`, you must provide the password either in this field or in the connection string of every application that uses this client DSN. If the server DSN is defined with `Authenticate=0`, then TimesTen ignores the value entered in this field, unless Access Control is enabled. In that case the Password is required.
3. In the **Timeout Interval** field, enter the interval time in seconds. You can enter any non-negative integer. A value of 0 indicates that Client/Server operations should not timeout. The default is 60 seconds. The maximum is 99,999 seconds.
4. Click **OK** to save the setup.

Deleting a server name

To delete a server name:

1. On the Windows Desktop on the Client machine, choose **Start > Settings > Control Panel**.
2. Double click **ODBC**. This opens the ODBC Data Source Administrator.

3. Click either **User DSN** or **System DSN**.
4. Select a TimesTen Client DSN and click **Configure**. This opens the TimesTen Client DSN Setup dialog.
5. Click **Servers**. This opens the TimesTen Logical Server List dialog.
6. Select a server name from the **TimesTen Servers** list.
7. Click **Delete**.

Accessing a remote data store on Windows

Example 2.1



In this example, the TimesTen Client machine is **client.mycompany.com**. The Client application is accessing the Server DSN on the remote server machine, `server.mycompany.com`. The logical name of the server is **ttserver_logical**.

Note: These examples reference the 32-bit sample DSNs. This is indicated by the extension **_32**. On 64-bit platforms, the sample DSNs are appended with **_64**.

1. On the server machine `server.mycompany.com`, use the **ttStatus** utility to verify that the TimesTen Server Daemon is running and to verify the port number it is listening on.
2. Using the procedure in “[Defining Server DSNs](#)” on page 43, verify that the Server DSN, **RunData70_32**, is defined as a System DSN on `server.mycompany.com`.
3. On the Client machine, `client.mycompany.com`, create a logical Server Name entry for the remote TimesTen Server. In the TimesTen Logical Server Name Setup dialog:
 - In the **Server Name** field, enter `ttserver_logical`.
 - In the **Network Address** field, enter `server.mycompany.com`.
 - In the **Network Port** field, enter 17002. This is the default port number for TimesTen Release 7.0. This value should correspond to the value displayed by **ttStatus** in Step 1.

See “[Creating and configuring a logical server name](#)” on page 45 for the procedure to open the TimesTen Server Name dialog and for more details.
4. On the Client machine, `client.mycompany.com`, create a Client DSN that corresponds to the remote Server DSN, **RunData_tt70_32**. In the TimesTen Client DSN Setup dialog, enter the following values:
 - In the **Client DSN** field, enter `RunDataCS_tt70_32`.
 - In the **Server Name or Network Address** field, enter `ttserver_logical`.
 - In the **Description** field, enter a description for the server. Entering data into this field is optional.
 - In the **Server DSN** field, enter `RunData70_32`.

5. Run the Client application from the machine `client.mycompany.com` using the Client DSN, **RunDataCS_tt70**. The example below uses the **ttIsqlCS** program installed with TimesTen Client.

```
ttIsqlCS connStr "DSN=RunDataCS_tt70_32"
```

Example 2.2 This example describes how to access a TimesTen Server that is listening on a port numbered other than the default port number.

Let us consider the Network Address of the TimesTen Server is `server.mycompany.com` and the Server is listening on Port 22222. The following methods can be used to connect to a Server DS:

1. Define the logical server name `logical_server` with `server.mycompany.com` as the Network Address and 22222 as the Network Port. Define a Client DSN with `logical_server` as the Server name, `Server_DSN` as the Server DSN. And execute the command:

```
ttIsqlCS -connStr "DSN=Client_DSN"
```

2. Alternatively, define the logical server name `logical_server` with `server.mycompany.com` as the Network Address and the default port number as the Network Port. Define a Client DSN with `logical_server` as the Server name, `Server_DSN` as the Server DSN. Overwrite the port number in the command:

```
ttIsqlCS -connStr "DSN=Client_DSN; TCP_Port=22222"
```

3. Alternatively, define the Server in the connection string. In this case you do not need to define a Client DSN, nor a logical server name.

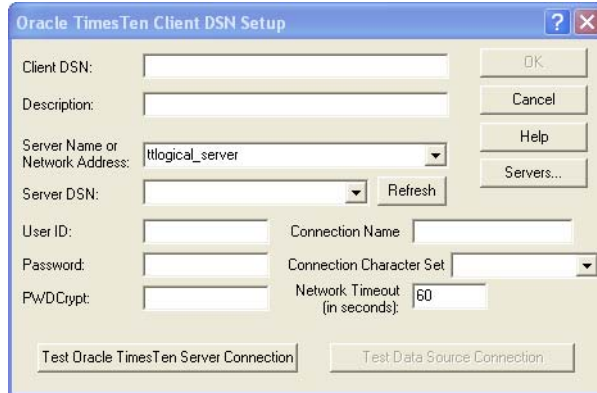
```
ttIsqlCS -connStr "TTC_Server=server.mycompany.com;  
TTC_Server_DSN=Server_DSN; TCP_Port=22222"
```

Testing connections

To test Client application connections to TimesTen data stores:

1. On the Windows Desktop, choose **Start > Settings > Control Panel**.
2. Double click **ODBC**. This opens the ODBC Data Source Administrator.
3. Click **User DSN** or **System DSN**.

4. Select the TimesTen Client DSN whose connection you want to test and click **Configure**. This opens the TimesTen Client DSN Setup dialog.



5. Click **Test TimesTen Server Connection** to test the connection to TimesTen Server.

The ODBC Data Source Administrator attempts to connect to TimesTen Server Daemon and displays messages to indicate if it was successful. During this test TimesTen Client verifies that:

- ODBC, Windows sockets and TimesTen Client are installed on the client machine.
 - The server specified in the **Server Name or Network Address** field of the TimesTen Client DSN Setup dialog is defined and the corresponding machine exists.
 - The TimesTen Server Daemon is running on the server machine.
6. Click **Test Data Source Connection** to test the connection to the Server DSN. The ODBC Data Source Administrator attempts to connect to the TimesTen Server DSN and displays messages to indicate whether it was successful.

During this test, TimesTen Client verifies that:

- The Server DSN specified in the **Server DSN** field is defined on the server machine.
- A Client application can connect to the Server DSN.

Creating and configuring Client DSNs on UNIX



On UNIX, you define logical server names by editing the TTCONNECT.INI file and you define client DSNs by editing the user ODBC.INI file (for user DNS) or the system ODBC.INI file (for system DSNs). For a description of user and system DSNs, see [“Data source names” on page 13](#).

This section includes the following topics:

- [Searching for a DSN](#)
- [Creating and configuring a logical server name](#)
- [Creating a Client DSN](#)

Searching for a DSN

When TimesTen looks for a specific DSN, it looks in the following locations in this order:

1. The file referenced by the ODBCINI environment variable, if it is set
2. The `.odbc.ini` file in the user's home directory, if the ODBCINI environment variable is not set
3. The file referenced by the SYSODBCINI environment variable, if it is set
4. One of the following, if the SYSODBCINI environment variable is not set:
 - The `InstallDir/info/sys.odbc.ini` file for a nonroot installation
 - The `/var/TimesTen/InstanceName/sys.odbc.ini` file for a root installation
5. The `/var/TimesTen/sys.odbc.ini` file, if the SYSODBCINI environment variable is not set

Creating and configuring a logical server name

You define logical server names in the `/var/TimesTen/instance/sys.ttconnect.ini` file or in a file named by the SYSTTCONNECTINI environment variable. This file is referred to as the TTCONNECT.INI file. The file contains a description, a network address and a port number.

The Network Address must be one of:

Type of Connection	Network Address
Local Client/Server connection that uses UNIX domain sockets	<code>ttLocalHost</code>
Local Client/Server connection that uses shared memory for inter-process communication	<code>ttShmHost</code>
Remote Client/Server connection	The name of the machine where the TimesTen Server is running. For example, <code>server.mycompany.com</code>

TimesTen searches for the logical Server in this order:

1. In the file specified by the SYSTTCONNECTINI environment variable, if it is set

2. In the `daemon_home_dir/sys.ttconnect.ini` file
3. Or in the default `/var/TimesTen/sys.ttconnect.ini` file

Examples

Example 2.3 Below is an example from a TTCONNECT.INI file that defines a logical server name, `ttserver_logical`, for a TimesTen Server daemon running on the machine `server.mycompany.com` and listening on port 17002. The instance name of the TimesTen installation is “tt70.”

```
[ttserver_logical]
Description=TimesTen Server 7.0
Network_Address=server.mycompany.com
TCP_Port=17002
```

Example 2.4 If both the client and server are on the same UNIX machine, applications using the TimesTen Client ODBC driver may improve performance by using UNIX domain sockets for communication.

The logical server name must also define the port number on which the TimesTen Server Daemon is listening so that multiple instances of the same version of TimesTen Server Daemon can be run on the same machine. To achieve this, the logical server name definition in TTCONNECT.INI file might look like:

```
[LocalHost_tt70]
Description=
  Local TimesTen Server TimesTen release 7.0 through domain
  sockets
Network_Address=ttLocalHost
TCP_PORT=17002
```

Example 2.5 If both the client and server are on the same machine, applications can use shared memory for inter-process communication. This may result in the best performance.

The logical server name must also define the port number on which the TimesTen Server Daemon is listening in order to make the initial connection. To achieve this, the logical server name definition in TTCONNECT.INI file might look like:

```
[ShmHost_tt70]
Description=
  Local TimesTen Server TimesTen release 7.0 through shared memory
Network_Address=ttShmHost
TCP_PORT=17002
```

Creating a Client DSN

In the ODBC Data Sources section of the ODBC.INI file, add an entry for the Client DSN. Each entry in this section lists the data source and the name of the ODBC driver that the data source uses. Use the following format for data source entries.

```
[ODBC Data Sources]
data-source-name=name-of-ODBC-driver
```

For example, to add the **RunDataCS_tt70_32** data source and associate it with the TimesTen Client ODBC driver, make the following entry in the ODBC Data Sources section of the ODBC.INI file.

```
[ODBC Data Sources]
RunDataCS_tt70_32=TimesTen Client 7.0
```

After the ODBC Data Sources section, add an entry to specify the connection attributes for each data source you have defined. Each data source listed in the ODBC Data Sources section of the ODBC.INI file requires a data source specification section.

The following is an example specification of the TimesTen Client example DSN **RunData_tt70_32**.

```
[RunDataCS_tt70_32]
TTC_Server=ttserver_logical
TTC_Server_DSN=RunData_tt70_32
```

For a description of the Client DSN attributes used in the ODBC.INI file, see [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*.

Accessing a remote data store on UNIX

Example 2.6



In this example, the TimesTen Client application machine is a 32-bit Solaris machine, `client.mycompany.com`. The Client application is accessing the Server DSN `RunData_tt70_32` on the remote server machine, another 32-bit Solaris machine, `server.mycompany.com`. The logical name of the server is `ttserver_logical`. The instance name of the TimesTen installation is “tt70.”

1. On the server machine `server.mycompany.com`, use the **ttStatus** utility to verify that the TimesTen Server is running and to verify the port number on which it is listening.
2. Verify that the Server DSN **RunData_tt70_32** exists in the system ODBC.INI file on `server.mycompany.com`.

There should be an entry in the ODBC.INI file as follows:

```
[RunData_tt70_32]
Driver=/opt/TimesTen/tt70/lib/libtten.so
```



```
DataStore=/var/TimesTen/tt70/server/RunData_tt70_32
```

3. Create a logical Server Name entry for the remote TimesTen Server in the TTCONNECT.INI file on `client.mycompany.com`.

```
[ttserver_logical]
# This value for TCP_Port should correspond to the
# value reported by ttStatus when verifying that the
# server is running
Network_Address=server.mycompany.com
TCP_Port=17002
```

See “[Creating and configuring Client DSNs on UNIX](#)” on page 51 for information on the creating a TTCONNECT.INI file.

4. On the Client machine, `client.mycompany.com`, create a Client DSN corresponding to the remote Server DSN, **RunData_tt70_32**.

There should be an entry in the ODBC.INI file as follows:

```
[RunDataCS_tt70_32]
TTC_SERVER=ttserver_logical
TTC_SERVER_DSN=RunData_tt70_32
```

See “[User and system DSNs](#)” on page 13 for information on the location of the proper ODBC.INI file.

5. Run the Client application from the machine `client.mycompany.com` using the Client DSN, **RunDataCS_tt70_32**. The example below uses the **ttIsql** program that is installed with TimesTen Client.

```
ttIsqlCS -connStr "DSN=RunDataCS_tt70_32"
```

Example 2.7 This example describes how to access a TimesTen Server that is listening on a port numbered other than the default port number.

Let us consider the Network Address of the TimesTen Server is `server.mycompany.com` and the Server is listening on Port 22222. The following methods can be used to connect to a Server DS:

1. Define the logical server name `logical_server` with `server.mycompany.com` as the Network Address and 22222 as the Network Port. Define a Client DSN with `logical_server` as the Server name, `Server_DSN` as the Server DSN. And execute the command:

```
ttIsqlCS -connStr "DSN=Client_DSN"
```

2. Alternatively, define the logical server name `logical_server` with `server.mycompany.com` as the Network Address and the default port number as the Network Port. Define a Client DSN with `logical_server` as the Server name, `Server_DSN` as the Server DSN. Overwrite the port number in the command:

```
ttIsqlCS -connStr "DSN=Client_DSN; TCP_Port=22222"
```

3. Alternatively, define the Server in the connection string. In this case you do not need to define a Client DSN, nor a logical server name.

```
ttIsqlCS -connStr "TTC_Server=server.mycompany.com;  
TTC_Server_DSN=Server_DSN; TCP_Port=22222"
```

Testing connections

To test Client application connections to TimesTen data stores:

1. Verify that the client machine can access the server machine.
2. Run `ping` from the client machine to see if a response is received from the server machine.
3. Verify that the TimesTen Server Daemon is running on the server machine.
 - Use `telnet` to connect to the port on which the TimesTen Server Daemon is listening. For example:

```
telnet server.mycompany.com 17002
```
 - If you successfully connect to the TimesTen Server Daemon, you will see a message similar to:

```
Connected to server.mycompany.com
```
 - If the server machine responds to a command, but TimesTen Server Daemon does not, the TimesTen Server Daemon may not be running. In the case of a failed connection, you will see a message similar to:

```
telnet: Unable to connect to remote host:  
Connection refused
```
 - Use the [ttStatus](#) utility on the server machine to determine the status and port number of the TimesTen Server. Generally, the TimesTen Server Daemon is started at installation time. If the TimesTen Server Daemon is not running, you must start it. For information on starting the TimesTen Server, see [“Modifying the TimesTen Server daemon options” on page 67](#).
4. Verify that the Client application can connect to the data store. If you cannot establish a connection to the data store, check that the TTCONNECT.INI file contains the correct information.
5. If the information in the TTCONNECT.INI file is correct, check that a Server DSN corresponding to the data store has been defined properly in the system ODBC.INI file on the machine where the data store resides and where the TimesTen Server Daemon is running.

Working with the TTCONNECT.INI file

TimesTen uses theTTCONNECT.INI file to define the names and attributes for servers and the mappings between logical server names and their network addresses. This information is stored on machine where the TimesTen Client is installed. By default, the TTCONNECT.INI file is `/var/TimesTen/sys.ttconnect.ini`.

To override the name and location of this file at runtime, set the `SYSTTCONNECTINI` environment variable to the name and location of the TTCONNECT.INI file before launching the TimesTen application.

Defining a server name on UNIX

You can define short-hand names for TimesTen Servers on UNIX in the TTCONNECT.INI file. The format of a TimesTen Server specification in the TTCONNECT.INI file is shown in [Table 2.1](#).

Table 2.1 TimesTen Server Format in the TTCONNECT.INI File

Components	Description
<code>[ServerName]</code>	Short name of the TimesTen you wish to define
<code>Description=description</code>	Description of the TimesTen Server
<code>Network_Address=network-address</code>	The DNS name, host name or IP address of the machine on which the TimesTen Server is running.
<code>TCP_Port=port-number</code>	The TCP/IP port number where the TimesTen Server is running. Default for TimesTen release 7.0 is 17002 for 32-bit platforms and 17003 for 64-bit platforms.

For example, the server specification for a remote TimesTen Server might appear as:

```
[ttserver]
Description=TimesTen Client/Server
Network_Address=server.company.com
TCP_Port=17002
```

For a local TimesTen Client/Server application that is using UNIX domain sockets, the network address must be defined as `ttLocalHost`. The server specification might appear as:

```
[LocalHost70]
Description=Shm TimesTen Client/Server
Network_Address=localhost
TCP_Port=17002
```

For a TimesTen Client/Server application that is using a shared memory segment for inter-process communication, the network address must be defined as `ttShmHost`. The server specification might appear as:

```
[ShmHost70]
Description=Shm TimesTen Client/Server
Network_Address=shmhost.company.com
TCP_Port=17002
```

Working with the Oracle TimesTen Data Manager Daemon

The Oracle TimesTen Data Manager daemon (Oracle TimesTen Data Manager service on Windows) starts when TimesTen is installed. If TimesTen was installed by the user `root`, the daemon also starts each time the operating system is booted. Otherwise, it starts after the start setup scripts have been run by `root`. The daemon operates continually in the background.

The TimesTen daemon performs the following functions:

- Manages shared memory access
- Coordinates process recovery
- Keeps management statistics on what data stores exist, which are in use, and which application processes are connected to which data stores
- Manages RAM policy
- Starts replication processes, the TimesTen Server daemon, the cache connect agent and the webserver

Application developers do not interact with the daemon directly. No application code runs in the daemon and application developers do not generally have to be concerned with it. Application programs that access TimesTen data stores communicate with the daemon transparently using TimesTen internal routines.

This chapter discusses interaction with the TimesTen daemon on various platforms. It includes the following topics:

- [Starting and stopping the Oracle TimesTen Data Manager Service on Windows](#)
- [Starting and stopping the daemon on UNIX](#)
- [Managing TimesTen daemon options](#)
- [Managing TimesTen Client/Server daemon options](#)
- [Modifying the TimesTen web server options](#)

Starting and stopping the Oracle TimesTen Data Manager Service on Windows



The Oracle TimesTen Data Manager service starts when you install the Oracle TimesTen Data Manager on your Windows system. To manually start and stop the Oracle TimesTen Data Manager service, you can use the **ttDaemonAdmin** utility with the `-start` or `-stop` option, or you can use Windows Administrative Tools as follows:

1. Open Administrative Tools:
On the Windows NT Desktop, choose **Start > Settings > Control Panel**.
On Windows 2000 and XP, choose **Start > Settings > Control Panel > Administrative Tools**.
2. Double-click **Services**. All currently available services are displayed.
3. Select **TimesTen Data Manager 7.0**, then click the appropriate button to stop or start the service.

Note: You must have administrative privileges to start and stop the TimesTen service.

Starting and stopping the daemon on UNIX



You must be root or the instance administrator to start and stop the TimesTen daemon.

For root installations, the TimesTen main daemon starts when TimesTen is installed and each time the operating system is booted. The daemon operates continually in the background.

For non-root installations, the instance administrator must manually start and stop the daemon, after each system reboot, unless the `setuproot` script has been run. To manually start and stop the TimesTen main daemon, you can use the **ttDaemonAdmin** utility with the `-start` or `-stop` option.

The following table shows the location of the TimesTen main daemon startup script by platform. These scripts exist only if the instance was installed by root or if the `setuproot` script has been run in a non-root instance.

To use the main daemon startup script to start the daemon, enter:

Platform	Command to start the TimesTen daemon
Solaris	<code># /etc/init.d/tt_TTinstance start</code>
HP-UX	<code># /sbin/init.d/tt_TTinstance start</code>

Linux	# /etc/rc.d/init.d/tt_ <i>TTinstance</i> start
AIX	# startsrc -s tt_ <i>TTinstance</i>
Tru64	# /sbin/init.d/tt_ <i>TTinstance</i> start

To use the main daemon startup script to stop the daemon, enter:

Platform	Command to stop the TimesTen daemon
Solaris	# /etc/init.d/tt_ <i>TTinstance</i> stop
HP-UX	# /sbin/init.d/tt_ <i>TTinstance</i> stop
Linux	# /etc/rc.d/init.d/tt_ <i>TTinstance</i> stop
AIX	# stopsrc -s tt_ <i>TTinstance</i>
Tru64	# /sbin/init.d/tt_ <i>TTinstance</i> stop

On an AIX system with a root installation, you can also determine the status of the daemon at any time by entering the command:

```
# lssrc -s TTinstance
```

Managing TimesTen daemon options

The `ttendaemon.options` file allows you to set and modify TimesTen daemon options. During installation, the installer sets some of these options to correspond to your responses to the installation prompts.

On Windows, the `ttendaemon.options` file is located in the directory:

```
install_dir\srv\info
```

On UNIX, the `ttendaemon.options` file is located in the directory:

```
/var/TimesTen/TTinstance/ (For root installations)
```

```
install_dir/info/ (For nonroot installations)
```

The features that the `ttendaemon.options` file controls are:

- The addresses on which the daemon listens
- The minimum and maximum number of TimesTen subdaemons that can exist for the TimesTen instance
- Whether or not the TimesTen Server daemon is started
- Whether or not you use shared memory segments for client/server inter-process communication
- The number of Server processes that are prespawnd on your system

- The location and size of support and user logs
- Backward compatibility
- The maximum number of users for a TimesTen instance where Access Control is enabled.
- Data access across NFS mounted systems (Linux only)
- The location of the Oracle Database installation (This option cannot be modified in this file.)

Note: To make changes to the `ttendaemon.options` file, you must first stop the TimesTen daemon and restart it after you have completed your changes.

For TimesTen Server options, it is only necessary to stop the server. It is not necessary to stop the TimesTen daemon.

The rest of this section includes the following topics:

- [Determining the daemon listening address](#)
- [Modifying informational messages](#)
- [Changing the allowable number of subdaemons](#)
- [Allowing data store access over NFS-mounted systems](#)
- [Enabling Linux large page support](#)
- [Other daemon options settings](#)

Determining the daemon listening address

By default, the TimesTen main daemon, all subdaemons and agents listen on a socket for requests, using any available address. All TimesTen utilities and agents use the loopback address to talk to the main daemon, and the main daemon uses the loopback address to talk to agents.

The `-listenaddr` entry in a separate line in the `ttendaemon.options` file tells the TimesTen daemons to listen on the specific address indicated in the value supplied. The address specified with this option can be either a host name or a numerical IP address.

The `-listenaddr` parameter exists for situations where a server has multiple network addresses and multiple network cards. In this case it is possible to limit the network addresses on which the TimesTen daemon is listening to a subset of the server's network addresses. This is done by making entries only for those addresses on which the daemon listens. These possibilities exist:

- Given a situation where a server has a “public” network address that is accessible both inside and outside the local network and a “private” address that is accessible only within the local network, adding a `-listenaddr` entry containing only the private address blocks all communications to TimesTen coming on the public address.

- By specifying only the localhost, the TimesTen main daemon can be cut off from all communications coming from outside the server and communicate only with local clients and subdaemons.

There is no relationship between TimesTen replication and the `-listenaddr` parameter and there is no requirement for enabling the `-listenaddr` parameter when replication is enabled. If replication is going to be used in an environment where `-listenaddr` is enabled, then the replication nodes need to know the allowable network addresses to use. However, if no `-listenaddr` parameter is enabled replication still works.

To explicitly specify the address on which the daemons should listen on a separate line in the `ttendaemon.options` file, enter:

```
-listenaddr address
```

For example, if you want to restrict the daemon to listen to just the loopback address, you say either:

```
-listenaddr 127.0.0.1
```

or

```
-listenaddr localhost
```

This means that only processes on the local machine can communicate with the daemon. Processes from other machines would be excluded, so you would not be able to replicate to or from other machines, or grant client access from other machines.

If you have multiple ethernet cards on different subnets, you can specify `-listenaddr` entries to control which machines can connect to the daemon.

You can enter up to four addresses on which to listen by specifying the option and a value on up to four separate lines in the `ttendaemon.options` file. In addition to the addresses you specify, TimesTen always listens on the loopback address.

Listening on IPv6


TimesTen supports the IPv6 protocol on Solaris and Linux systems. By default, TimesTen uses the IPv4 protocol. To enable the daemon to listen on IPv6, you must enter on a separate line in the `ttendaemon.options` file:

```
-enableIPv6
```

and

```
-listenaddr6 address
```

You can specify just the `-listenaddr6` option to enable IPv6. But, if you specify `-enableIPv6` without specifying the `-listenaddr6` option, IPv6 is not enabled. Specifying `-enableIPv6` option with `-listenaddr`, adds the IPv6 loopback interface to the list.



The address specified with this option can be either a host name or a numerical IP address. See “[Determining the daemon listening address](#)” on page 62 for specifics on the `-listenaddr` option

If one or more `-listenaddr` options are provided, the daemons listen on the specified IPv4 interfaces, with the IPv4 loopback address being added to the list if not specified. If only `-enableIPv6` is specified, the daemons listen on both the IPv4 ANY interface and the IPv6 ANY interface.

You can specify both `-listenaddr` and `-listenaddr6` options. If you specify one or more `-listenaddr6` options, the daemons listen on the specified IPv4 or IPv6 interfaces, with both the IPv4 and IPv6 loopback interfaces being added if not specified. If the name resolver returns multiple IPv4 and/or IPv6 addresses for a name, the daemons listen on all of the names.

Modifying informational messages

As the daemon operates, it generates error, warning and informational messages. These messages may be useful for TimesTen system administration and for debugging applications.

By default, informational messages are stored in:

- A user error log that contains information you may need to see. Generally, these messages contain information on actions you may need to take.
- A support log containing everything in the user error log plus information of use by TimesTen Customer Support.

The following options specify the locations and size of the support and user logs, as well as the number of files to keep stored on your system.

Option	Description
<code>-supportlog path</code> <code>-f path</code>	Specifies the location for the support log file. The default file is <code>daemon_home/ttmesg.log</code>
<code>-maxsupportlogfiles num</code>	The TimesTen main daemon automatically rotates the files once they get to a specific size. This option specifies the number of support log files to keep. The default is 10.
<code>-maxsupportlogsize nBytes</code>	Specifies the maximum size of the support log file. The default is 1MB.

Option	Description
<code>-userlog path</code>	Specifies the location for the user log file. The default file is <code>daemon_home/tterrors.log</code> . You may specify <code>syslog</code> on UNIX systems as the path or the Event Log on Windows, in which case the output is sent to the system <code>syslog</code> or Event Log.
<code>-maxuserlogfiles num</code>	The TimesTen main daemon automatically rotates the files once they get to a specific size. This option specifies the number of user log files to keep. The default is 10.
<code>-maxuserlogsize nBytes</code>	Specifies maximum size of the user log. Default is 1MB.
<code>-showdate</code>	On UNIX systems only, indicates that the date should be prepended to all messages



On Windows, if you have specified the Event Log as the location for your log messages, to view them follow these steps:

1. Open the **Event Viewer** window on your Windows Desktop.
2. From the Log menu, choose **Application**.

The window changes to display only log messages generated by applications. Any messages with the word “TimesTen” in the “Source” column were generated by the Oracle TimesTen Data Manager service.

3. To view any TimesTen message, double-click the message summary.

The message window is displayed. You can view additional messages by clicking **Next / Previous** or the up / down arrows, depending on your version of Windows.

Note: You can also view messages using the **ttDaemonLog** utility.



On UNIX systems, daemon messages are routed in a variety of ways, including recording them to a file. These files can grow to be quite large. You should prune them periodically to conserve disk space.

To specify the `syslog` facility used to log TimesTen Daemon and subdaemon messages, on a separate line of the `ttendaemon.options` file add:

`-facility name`

Possible `name` values are: `auth`, `cron`, `daemon`, `local0-local7`, `lpr`, `mail`, `news`, `user`, or `uucp`.

To turn off detailed log messages, add a “#” before `-verbose` of the `ttendaemon.options` file.

Changing the allowable number of subdaemons

TimesTen uses subdaemons to manage data stores. The main TimesTen daemon spawns subdaemons dynamically as they are needed. You can manually specify a range of subdaemons that the daemon may spawn, by specifying a minimum and maximum.

At any point in time, one subdaemon is potentially needed for TimesTen process recovery for each failed application process that is being recovered at that time.

By default, the maximum number of subdaemons is 50.

By default, TimesTen spawns a minimum of 4 subdaemons. However, you can change these settings by assigning new values to the `-minsubs` and `-maxsubs` options in the `ttendaemon.options` file.



Allowing data store access over NFS-mounted systems

By default, TimesTen systems cannot access data across NFS-mounted systems. On Linux x86 32-bit and 64-bit systems, you can access checkpoint and log files on NFS-mounted systems.

To enable data access on NFS-mounted systems, on a separate line of the `ttendaemon.options` file, add:

```
-allowNetworkFiles
```

Note: TimesTen does not support the storage of trace files or user and support logs across NFS-mounted systems.



Enabling Linux large page support

To enable Linux large page support on TimesTen, on a separate line of the `ttendaemon.options` file, add:

```
-linuxLargePageAlignment Size_in_MB
```

The *Size_in_MB* is the `Hugepagesize` value in `/proc/meminfo`, specified in MB instead of KB.



Other daemon options settings

To limit the number of users that have access to a TimesTen instance when Access control is enabled on a TimesTen instance, on a separate line of the `ttendaemon.options` file, add:

`-maxusers num`

To view the location of the Oracle Database installation (ORACLE_HOME), read the value for the option:

`-oracle_home path`

Note: The only way to change the location of ORACLE_HOME after the installation of TimesTen is to use the utility **ttmodinstall**.

Managing TimesTen Client/Server daemon options

This section includes the following topics:

- [Modifying the TimesTen Server daemon options](#)
- [Controlling the TimesTen Server daemon](#)
- [Prespawning TimesTen Server processes](#)
- [Using shared memory for Client/Server IPC](#)
- [Controlling the TimesTen Server log messages](#)
- [Communicating with older versions of TimesTen](#)

Modifying the TimesTen Server daemon options

The TimesTen Server is a subdaemon of the TimesTen daemon that operates continually in the background. To modify the TimesTen Server daemon options, you must:

1. Stop the TimesTen Server
2. Modify the options in the `ttendaemon.options` file as described in the following sections
3. Restart the TimesTen Server.

Controlling the TimesTen Server daemon

The `-server portno` entry in a separate line in the `ttendaemon.options` file tells the TimesTen daemon to start the TimesTen Server daemon and what port to use. The `portno` is the port number on which the server will listen.

If the TimesTen Server is installed, you can enable or disable the TimesTen Server daemon by:

- To enable the Server daemon, remove the comment symbol '#' in front of the `-server portno` entry.
- To disable the Server daemon, add a comment symbol '#' in front of the `-server portno` entry.

Prespawning TimesTen Server processes

Each TimesTen Client connection requires one server process. By default, a server process is spawned when a client requests a connection.

You can prespawn a pool of reserve server processes, making them immediately available for a client connection, thus improving client/server connection performance.

The `-serverpool number` entry in a separate line in the `ttendaemon.options` file on the Server machine tells the TimesTen Server daemon create `number` processes. If this option is not specified, no processes are pre-spawned and kept in the reserve pool.

When a new connection is requested, if there are no items in the server pool, a new process is spawned, as long as you have not met the operating system limit.

If you request more process than allowed by your operating system, a warning is returned. Regardless of the number of processes requested, an error does not occur unless a client requests a connection when no more are available on the system, even if there are no processes remaining in the reserve pool.

Note: These changes to the TimesTen Server daemon do not occur until the TimesTen daemon is restarted.

Specifying multiple connections to the TimesTen Server

By default, TimesTen creates only one connection to a Server per child process. You can set multiple connects to a single TimesTen Server, either by using the [Server connection attributes](#) described in the *Oracle TimesTen In-Memory Database API Reference Guide* or by setting the TimesTen daemon options described in this section. These options allow you to set the number of connections to a TimesTen server, the number of servers for each DSN and the size of each connection to the server.

Note: In the case that you have set both the Server connection attributes and these daemon options, the value of the Server connection attributes takes precedence.

Configuring the maximum number of client connections per child server process

To run a child server process in multithreaded mode so that a single server process can service multiple client connections to a data store, add the following line to the `ttendaemon.options` file:

`-maxConnsPerServer NumberOfClientConnections`

The possible values of *NumberOfClientConnections* range from 1 to 2047, inclusive. The default value is 1, which indicates that the child server process runs in multi-process mode and, therefore, can service only one client connection.

Configuring the desired number of child server processes spawned for a server DSN

To specify the desired number of child server processes to be spawned for a particular server DSN, add the following line to the `ttendaemon.options` file:

`-serversPerDSN NumberOfChildServerProcesses`

The possible values of *NumberOfChildServerProcesses* range from 1 to 2047, inclusive. The default value is 1.

Client connections to a particular server DSN are evenly distributed in round-robin fashion to the child server processes that are spawned and assigned to the DSN. The number of child server processes assigned to the server DSN is greater than *NumberOfChildServerProcesses* if the number of client connections to the DSN is greater than the maximum number of client connections per child server process multiplied by the desired number of child server processes spawned for a server DSN.

Configuring the Thread Stack Size of the Child Server Processes

To set the size of the child server process thread stack for each client connection, add the following line to the `ttendaemon.options` file:

`-serverStackSize ThreadStackSize`

ThreadStackSize is specified in KB. The default is 128 KB on 32-bit systems and 256 KB on 64-bit systems. The *ThreadStackSize* setting is ignored if the maximum number of client connections per child server process is 1 because the sole client connection will be serviced by the main thread of the child server process.

Note: These changes to the TimesTen Server daemon do not occur until the TimesTen daemon is restarted.

Using shared memory for Client/Server IPC

By default, TimesTen uses TCP/IP communication between applications linked with the TimesTen Client driver and the TimesTen Server.

Where the client application resides on the same machine as the TimesTen Server, you can alternatively use shared memory for the inter-process communication (IPC).

This can be useful for performance purposes or to allow 32-bit client applications to communicate with a 64-bit data store on the server. Before using shared memory as IPC verify that you have configured your system correctly. See "Installation prerequisites" in Chapter 2 of the *Oracle TimesTen In-Memory Database Installation Guide*.

The `-serverShmIpc` entry in a separate line in the `ttendaemon.options` file tells the TimesTen Server daemon to accept a client connection that intends to use a shared memory segment for IPC.

If this entry is missing, add this line to the `ttendaemon.options` file to start the TimesTen Server daemon with shared memory IPC capability when the TimesTen daemon is restarted.

If the entry exists, add the `#` symbol before the line in the `ttendaemon.options` file to comment it out. The TimesTen Server is no longer started with shared memory IPC capability when the TimesTen daemon starts.

Note: TimesTen supports a maximum of 16 different instances of the shared memory IPC-enabled server. If an application tries to connect to more than 16 different shared memory segments it receives an ODBC error.

Managing the size of the shared memory segment

The `-serverShmSize` *size* entry in a separate line in the `ttendaemon.options` file tells the TimesTen Server daemon to create a shared memory segment of the specified size in MB.

If this entry is missing, the TimesTen Server daemon creates a shared memory segment of 64MB.

An appropriate value for the shared memory segment depends on:

- The expected number of concurrent client/server connections to all data stores that belong to an instance of the TimesTen Server.
- The number of concurrent allocated statements within each such connection.
- The amount of data being transmitted for a query.

Some guidelines for determining the size of the shared memory segment include:

- TimesTen needs 1 MB of memory for internal use.
- Each connection needs a fixed block of 16 KB.
- Each statement starts with a block of 16 KB for the IPC. But, this size is increased or decreased depending upon the size of the data being transmitted for a query. TimesTen increments the statement buffer size by doubling it and decreases it by halving it.

For example, if the user application anticipates a max of 100 simultaneous shared-memory-enabled client/server connections, and if each connection is

anticipated to have a maximum of 50 statements, and the largest query returns 128 KB of data, use this formula to configure the `serverShmSize`:

$$\begin{aligned}\text{serverShmSize} &= 1 \text{ MB} + (100 * 16) \text{ KB} + (100 * 50 * 128) \text{ KB} \\ &= 1 \text{ MB} + 2 \text{ MB} + 625 \text{ MB} = 628 \text{ MB}\end{aligned}$$

This is the most memory required for this example. The entire memory segment would be used only if all 100 connections have 50 statements each and each statement has a query that returns 128 KB of data in a row of the result.

In this example, if you configured the `serverShmSize` to 128 MB, either a new shared-memory-enabled client/server connection is refused by the TimesTen Server or a query may fail due to lack of resources within the shared memory segment.

Changing the size of the shared memory segment

Once configured, to change the value of the shared memory segment you must stop the TimesTen Server. Stopping the server detaches all existing client/server connections to any data store that is associated with that instance of the TimesTen Server. The steps for modifying the value of the `-serverShmSize` option are:

1. Modify the value of `-serverShmSize` in the `ttendaemon.options` file.
2. Use the **ttDaemonAdmin** utility to restart the TimesTen Server.

Controlling the TimesTen Server log messages

The `-noserverlog` entry in a separate line in the `ttendaemon.options` file tells the TimesTen daemon to turn off logging of connects and disconnects from the client applications.

If the TimesTen Server is installed, you can enable or disable logging of connect and disconnect messages by:

- To enable logging, add a comment symbol '#' before the `-noserverlog` entry.
- To disable logging, remove the comment symbol '#' before the `-noserverlog` entry.

Communicating with older versions of TimesTen

In TimesTen 7.0, the security of the communication protocols between the client and server and between the replication daemons is significantly improved over previous versions. Because of this, older versions of TimesTen using less secure protocols are not normally allowed to communicate with TimesTen 7.0.

However, in some cases, such as when performing an online upgrade of your data stores, you may need to allow older, less secure versions of TimesTen to communicate with TimesTen 7.0.

To allow client/server and replication communication between TimesTen 7.0 and previous versions of TimesTen, on a separate line in the `ttendaemon.options` file, add:

```
-insecure-backwards-compat
```

Note: By using the `-insecure-backwards-compat` option, you are reducing the security of your TimesTen installation. You should only use the option when absolutely necessary, and discontinue its use when you no longer need it.

Modifying the TimesTen web server options

The TimesTen web server is a subdaemon of the TimesTen daemon that operates continually in the background. The TimesTen web server allows users to use the Cache Administrator, a web-based interface that allows you to work with cache groups. For more information on cache groups, see the *TimesTen Cache Connect to Oracle Guide*. If the TimesTen web server is specified in the `ttendaemon.options` file, then the TimesTen web server starts and stops with the TimesTen daemon.

To change whether the TimesTen web server is started by the daemon, perform the following tasks:

1. Stop the TimesTen daemon.
2. Modify the `ttendaemon.options` file as described in "[Controlling the TimesTen web server](#)".
3. Restart the TimesTen daemon.

Controlling the TimesTen web server

The webserver is available if you have installed the Cache Connect to Oracle option of TimesTen.

The existence of a `-webserver` entry in a separate line in the `ttendaemon.options` file tells the TimesTen daemon to start the TimesTen web server.

If the TimesTen web server is installed, you can enable or disable the TimesTen web server by:

- To start the web server, remove the comment symbol '#' in front of the `-webserver` entry.
- To stop the web server, add a comment symbol '#' in front of the `-webserver` entry.

Note: If the Cache Connect to Oracle option is not installed setting this option displays a warning if you execute the ttStatus utility. These changes do not take effect until the TimesTen daemon is restarted.

You can also start and stop the web server by using the **ttDaemonAdmin** utility.

Globalization Support

This chapter describes TimesTen globalization support features. It includes the following topics:

- [Overview of globalization support features](#)
- [Choosing a database character set](#)
- [Length semantics and data storage](#)
- [Connection character set](#)
- [Linguistic sorts](#)
- [SQL string and character functions](#)
- [Setting globalization support attributes](#)
- [Globalization support during migration](#)
- [Supported Character Sets](#)
- [Supported Linguistic Sorts](#)

Overview of globalization support features

TimesTen globalization support includes the following features:

- Character set support
You must choose a database character set when you create a data store. See [“Supported Character Sets” on page 86](#) for a list of supported character sets. You can also choose a connection character set for a session. See [“Connection character set” on page 78](#).
- Length semantics
You can specify byte semantics or character semantics for defining the storage measurement of character data types. See [“Length semantics and data storage” on page 77](#).
- Linguistic sorts and indexes. You can sort data based on linguistic rules. See [“Linguistic sorts” on page 79](#). You can use linguistic indexes to improve performance of linguistic sorts. See [“Using linguistic indexes” on page 81](#).
- SQL string and character functions

TimesTen provides SQL functions that return information about character strings. TimesTen also provides SQL functions that return a character from an encoded value. See [“SQL string and character functions” on page 82](#).

Note: This release of TimesTen does not support session language and territory

Choosing a database character set

TimesTen uses the database character set to define the encoding of data stored in character data types such as CHAR and VARCHAR2.

Use the [DatabaseCharacterSet](#) data store attribute to specify the database character set during data store creation. You cannot alter the database character set after data store creation, and there is no default value for [DatabaseCharacterSet](#). See [“Supported Character Sets” on page 86](#) for a list of supported character sets.

Consider the following questions when you choose a character set for a data store:

- What languages does the data store need to support now and in the future?
- Is the character set available on the operating system?
- What character sets are used on clients?
- How well does the application handle the character set?
- What are the performance implications of the character set?

If you are using Cache Connect to Oracle to cache Oracle tables, you must create the data store with the same database character set as the Oracle database.

This section includes the following topics:

- [Character sets and languages](#)
- [Client operating system and application compatibility](#)
- [Performance and storage implications](#)
- [Character sets and replication](#)

Character sets and languages

Choosing a database character set determines what languages can be represented in the database.

A group of characters (for example, alphabetic characters, ideographs, symbols, punctuation marks, and control characters) can be encoded as a character set. An encoded character set assigns unique numeric codes to each character in the character repertoire. The numeric codes are called code points or encoded values.

Character sets can be single-byte or multibyte. Single-byte 7-bit encoding schemes can define up to 128 characters and normally support just one language.

Single-byte 8-bit encoding schemes can define up to 256 characters and often support a group of related languages. Multibyte encoding schemes are needed to support ideographic scripts used in Asian languages like Chinese or Japanese because these languages use thousands of characters. These encoding schemes use either a fixed number or a variable number of bytes to represent each character. Unicode is a universal encoded character set that enables information from any language to be stored using a single character set. Unicode provides a unique code value for every character, regardless of the platform, program, or language.

Client operating system and application compatibility

The database character set is independent of the operating system. On an English operating system, you can create and run a database with a Japanese character set. However, when an application in the client operating system accesses the database, the client operating system must be able to support the database character set with appropriate fonts and input methods. For example, you cannot insert or retrieve Japanese data on the English Windows operating system without first installing a Japanese font and input method. Another way to insert and retrieve Japanese data is to use a Japanese operating system remotely to access the database server.

If all client applications use the same character set, then that character set is usually the best choice for the database character set. When client applications use different character sets, the database character set should be a superset of all the application character sets. This ensures that every character is represented when converting from an application character set to the database character set.

Performance and storage implications

For best performance, choose a character set that avoids character set conversion and uses the most efficient encoding for the languages desired. Single-byte character sets result in better performance than multibyte character sets, and they also are the most efficient in terms of space requirements. However, single-byte character sets limit how many languages you can support.

Character sets and replication

All data stores in a replication scheme must have the same database character set. No character set conversion occurs during replication.

Length semantics and data storage

In single-byte character sets, the number of bytes and the number of characters in a string are the same. In multibyte character sets, a character or code point consists of one or more bytes. Calculating the number of characters based on byte lengths can be difficult in a variable-width character set. Calculating column

lengths in bytes is called **byte semantics**, while measuring column lengths in characters is called **character semantics**.

Character semantics is useful for defining the storage requirements for multibyte strings of varying widths. For example, in a Unicode database (AL32UTF8), suppose that you need to define a VARCHAR2 column that can store up to five Chinese characters together with five English characters. Using byte semantics, this column requires 15 bytes for the Chinese characters, which are three bytes long, and 5 bytes for the English characters, which are one byte long, for a total of 20 bytes. Using character semantics, the column requires 10 characters.

The expressions in the following list use byte semantics. Note the BYTE qualifier in the CHAR and VARCHAR2 expressions.

- CHAR (5 BYTE)
- VARCHAR2(20 BYTE)

The expressions in the following list use character semantics. Note the CHAR qualifier in the VARCHAR2 expression.

- VARCHAR2(20 CHAR)
- SUBSTR(*string*, 1, 20)

By default, the CHAR and VARCHAR2 character data types are specified in bytes, not characters. Therefore, the specification CHAR(20) in a table definition allows 20 bytes for storing character data.

The **NLS_LENGTH_SEMANTICS** general connection attribute determines whether a new column of character data type uses byte or character semantics. It enables you to create CHAR and VARCHAR2 columns using either byte-length or character-length semantics without having to add the explicit qualifier. NCHAR and NVARCHAR2 columns are always character-based. Existing columns are not affected.

The default value for **NLS_LENGTH_SEMANTICS** is BYTE. Specifying the BYTE or CHAR qualifier in a data type expression overrides the **NLS_LENGTH_SEMANTICS** value.

Connection character set

The database character set determines the encoding of CHAR and VARCHAR2 character data types. The connection character set is used to describe the encoding of the incoming and outgoing application data, so that TimesTen can perform the necessary character set conversion between the application and the database. For example, this allows a non-Unicode application to communicate with a Unicode (AL32UTF8) database.

The **ConnectionCharacterSet** general connection attribute sets the character encoding for the connection, which can be different than the database character set. The connection uses the connection character set for information that passes

through the connection, such as parameters, SQL query text, results and error messages. Choose a connection character set that matches the application environment or the character set of your data source.

Best performance results when the connection character set and the database character set are the same because no conversion occurs. When the connection character set and the database character set are different, data conversion is performed in the ODBC layer. Characters that cannot be converted to the target character set are changed to replacement characters.

The default connection character set is US7ASCII. This setting applies to both direct and client connections.

Linguistic sorts

Different languages have different sorting rules. Text is conventionally sorted inside a database according to the binary codes used to encode the characters. Typically, this does not produce a sort order that is linguistically meaningful. A linguistic sort handles the complex sorting requirements of different languages and cultures. It enables text in character data types (CHAR, VARCHAR2, NCHAR, and NVARCHAR2) to be sorted according to specific linguistic conventions.

A linguistic sort operates by replacing characters with numeric values that reflect each character's proper linguistic order. TimesTen offers two kinds of linguistic sorts: monolingual and multilingual.

This section includes the following topics:

- [Monolingual linguistic sorts](#)
- [Multilingual linguistic sorts](#)
- [Case-insensitive and accent-insensitive linguistic sorts](#)
- [Performing a linguistic sort](#)
- [Using linguistic indexes](#)

Monolingual linguistic sorts

TimesTen compares character strings in two steps for monolingual sorts. The first step compares the major value of the entire string from a table of major values. Usually, letters with the same appearance have the same major value. The second step compares the minor value from a table of minor values. The major and minor values are defined by TimesTen. TimesTen defines letters with accent and case differences as having the same major value but different minor values.

Monolingual linguistic sorting is available only for single-byte and Unicode database character sets. If a monolingual linguistic sort is specified when the database character set is non-Unicode multibyte, then the default sort order is the binary sort order of the database character set.

For a list of supported sorts, see [“Supported Linguistic Sorts” on page 90](#).

Multilingual linguistic sorts

TimesTen provides multilingual linguistic sorts so that you can sort data for multiple languages in one sort. Multilingual linguistic sort is based on the ISO/OEC 14651 - International String Ordering and the Unicode Collation algorithm standards. This framework enables the database to handle languages that have complex sorting rules (such as those in Asian languages), as well as providing linguistic support for databases with multilingual data.

In addition, multilingual sorts can handle canonical equivalence and supplementary characters. Canonical equivalence is a basic equivalence between characters or sequences of characters. For example, ç is equivalent to the combination of c and ,.

For example, TimesTen supports a monolingual French sort (FRENCH), but you can specify a multilingual French sort (FRENCH_M). _M represents the ISO 14651 standard for multilingual sorting. The sorting order is based on the GENERIC_M sorting order and can sort accents from right to left. TimesTen recommends using a multilingual linguistic sort if the tables contain multilingual data. If the tables contain only French, then a monolingual French sort may have better performance because it uses less memory. It uses less memory because fewer characters are defined in a monolingual French sort than in a multilingual French sort. There is a trade-off between the scope and the performance of a sort.

For a list of supported multilingual sorts, see [“Supported Linguistic Sorts” on page 90](#).

Case-insensitive and accent-insensitive linguistic sorts

Operations inside a database are sensitive to the case and the accents of the characters. Sometimes you might need to perform case-insensitive or accent-insensitive comparisons.

To specify a case-insensitive or accent-insensitive sort:

- Append _CI to a TimesTen sort name for a case-insensitive sort. For example:
 - BINARY_CI: accent-sensitive and case-insensitive binary sort
 - GENERIC_M_CI: accent-sensitive and case-insensitive GENERIC_M sort
- Append _AI to a TimesTen sort name for an accent-insensitive and case-insensitive sort. For example:
 - BINARY_AI: accent-insensitive and case-insensitive binary sort
 - FRENCH_M_AI: accent-insensitive and case-insensitive FRENCH_M sort

Performing a linguistic sort

The **NLS_SORT** data store connection attribute indicates which collating sequence to use for linguistic comparisons. The **NLS_SORT** value affects the SQL string comparison operators and the ORDER BY clause.

You can use the **ALTER SESSION** statement to change the value of **NLS_SORT**.

Example 4.1 ALTER SESSION SET NLS_SORT=SWEDISH;
SELECT product_name
FROM product
ORDER BY product_name;

```
PRODUCT NAME
-----
aerial
Antenne
Lcd
ächzen
Ähre
```

You can also override the **NLS_SORT** setting by using the **NLSSORT** SQL function to perform a linguistic sort.

Example 4.2 SELECT * FROM test ORDER BY NLSSORT(name, 'NLS_SORT=SPANISH');

For more extensive examples of using **NLSSORT**, see "**NLSSORT**" in *Oracle TimesTen In-Memory Database SQL Reference Guide*.

Using linguistic indexes

You can create a linguistic index to achieve better performance during linguistic comparisons. A linguistic index requires storage for the sort key values.

To create a linguistic index, use a statement similar to the following:

```
CREATE INDEX german_index ON employees
(NLSSORT(employee_id, 'NLS_SORT=GERMAN'));
```

The optimizer chooses the appropriate index based on the values for **NLSSORT** and **NLS_SORT**.

You must create multiple linguistic indexes if you want more than one linguistic sort on a column. For example, if you want both GERMAN and GERMAN_CI sorts against the same column, create two linguistic indexes.

For more information, see "**CREATE INDEX**" in *Oracle TimesTen In-Memory Database SQL Reference Guide*.

SQL string and character functions

The following table summarizes SQL functions that operate on character strings:

SQL Function	Description
ASCIISTR	Takes as its argument either a string or an expression that resolves to a string in any character set. It returns the ASCII version of the string in the database character set. Non-ASCII characters are converted to Unicode escapes.
INSTR INSTRB INSTR4	Determines the first position, if any, at which one string occurs within another string. INSTRB uses bytes instead of characters. INSTR4 uses UCS4 code points.
LENGTH LENGTHB LENGTH4	Returns the length of a character string in an expression as number of characters. LENGTHB uses bytes instead of characters. LENGTH4 uses UCS4 code points.
LOWER and UPPER	The LOWER function converts expressions of type CHAR, NCHAR, VARCHAR2 or NVARCHAR2 to lowercase. The UPPER function converts expressions of type CHAR, NCHAR, VARCHAR2 or NVARCHAR2 to uppercase. Character semantics is supported for CHAR and VARCHAR2 types. The data type of the result is the same as the data type of the expression.
RTRIM	Removes trailing spaces from CHAR, VARCHAR2, NCHAR or NVARCHAR2 strings.
SUBSTR SUBSTRB SUBSTR4	Returns a VARCHAR2 or NVARCHAR2 string that represents a substring of a CHAR or NCHAR string. The returned substring is a specified number of characters, beginning from a designated starting point. SUBSTRB uses bytes instead of characters. SUBSTR4 uses UCS4 code points.
UNISTR	Takes as its argument a string that resolves to data of type NVARCHAR2. It returns the value in UTF-16 format. Unicode escapes are supported.

The following functions return characters:

- **CHR**: Returns the character with the specified binary value in the database character set.
- **NCHR**: Returns the character with the specified Unicode value.

See "Expressions" in *Oracle TimesTen In-Memory Database SQL Reference Guide* for more information including examples.

Setting globalization support attributes

The globalization support attributes are summarized in the following table:

Parameter	Description
DatabaseCharacterSet	Indicates the character encoding used by a data store.
ConnectionCharacterSet	Determines the character encoding for the connection, which may be different from the database character set.
NLS_SORT	Indicates the collating sequence to use for linguistic comparisons.
NLS_LENGTH_SEMANTICS	Sets the default length semantics.
NLS_NCHAR_CONV_EXCP	Determines whether an error is reported when there is data loss during an implicit or explicit data type conversion between NCHAR/NVARCHAR2 data and CHAR/VARCHAR2 data.

DatabaseCharacterSet must be set during data store creation. There is no default. See "Choosing a database character set" on page 76.

The rest of the attributes are set during connection to a data store. For more information about **ConnectionCharacterSet**, see "Connection character set" on page 78.

You can use the **ALTER SESSION** statement to change the following attributes during a session:

- NLS_SORT
- NLS_LENGTH_SEMANTICS
- NLS_NCHAR_CONV_EXCP

For more information, see "ALTER SESSION" in the *Oracle TimesTen In-Memory Database SQL Reference Guide* and "Data Store Attributes" in *Oracle TimesTen In-Memory Database API Reference Guide*.

Backward compatibility using TIMESTEN8

TIMESTEN8 is a restricted database character set that specifies behavior from TimesTen releases before 7.0. It is supported for backward compatibility only.

TIMESTEN8 has the following restrictions:

- There is no support for character set conversion of any kind. This includes:
 - Conversions between the application and the database. If **DatabaseCharacterSet** is TIMESTEN8, then **ConnectionCharacterSet** must also be TIMESTEN8.
 - Conversions between CHAR VARCHAR2 data and NCHAR/NVARCHAR2 data.
- Sorting for CHAR and VARCHAR data types is limited to binary ordering. `NLS_SORT=BINARY` is the only sort allowed.
- TIMESTEN8 is not supported in Cache Connect to Oracle.

During database creation, customers should select the database character set matching the actual encoding of the data being stored in CHAR and VARCHAR2 columns whenever possible. Select TIMESTEN8 only when backwards compatibility to existing TimesTen data is required.

Globalization support during migration

The **ttMigrate** utility saves one or more migrate objects from a TimesTen data store into a binary data file or restores the objects from the binary data files into a TimesTen data store. Migrate objects include tables, cache group definitions, views and sequences.

This section includes the following topics:

- [Object migration and character sets](#)
- [Migration and length semantics](#)
- [Migrating linguistic indexes](#)
- [Migrating cache group tables](#)

See also "Copying, migrating, backing up and restoring a data store" on page 32 of this guide and the description of **ttMigrate** in *Oracle TimesTen In-Memory Database API Reference Guide*.

Object migration and character sets

ttMigrate tags each object it saves with the object's character set. By default, **ttMigrate** stores object data in the database character set, but you can choose a

different character set by using the `-saveAsCharset` option. You can specify this option in create mode (`-c`) or append mode (`-a`).

When you use **ttMigrate** to restore an object, its data is implicitly converted to the database character set of the target data store if needed. Character set conversion can result in data loss if the database character set of the target data store cannot represent all of the data that it receives.

If you know that the data is encoded in the database character set of the target data store, you can use the `-noCharsetConversion` option. This option can be specified only in restore mode (`-r`). If you use the `-noCharsetConversion` option, **ttMigrate** treats the data as if it is in the database character set of the target data store.

When you restore untagged character data from a data store that was created before release 7.0 into a data store from release 7.0 and later, **ttMigrate** treats the data as if it is in the database character set of the target data store.

ttMigrate issues a warning whenever there is an implicit or explicit character set conversion while saving or restoring data.

Migration and length semantics

ttMigrate saves length semantic information about CHAR and VARCHAR2 columns. It restores the length semantic information when restoring objects into data stores created in TimesTen release 7.0 or later.

When objects are migrated back into a TimesTen release before 7.0, columns with character semantics are converted to byte semantics and the column length is adjusted to match the byte length of the original columns.

When objects are migrated from a release before 7.0 to release 7.0 and later, byte semantics is used.

Migrating linguistic indexes

ttMigrate supports migration of linguistic indexes into TimesTen releases that support them. When migrating back to a TimesTen release before 7.0, **ttMigrate** issues a warning indicating that the linguistic indexes cannot be restored. Migration of the table proceeds without the linguistic indexes.

Migrating cache group tables

You cannot restore cache group tables containing NCHAR/NVARCHAR2 columns to a release before 7.0. Releases before 7.0 do not allow these data types in cache group tables.

Supported Character Sets

The tables in this section describe the character sets supported in TimesTen:

Asian Character Sets

Name	Description
JA16EUC	EUC 24-bit Japanese
JA16EUCTILDE	The same as JA16EUC except for the way that the wave dash and the tilde are mapped to and from Unicode
JA16SJIS	Shift-JIS 16-bit Japanese
JA16SJISTILDE	The same as JA16SJIS except for the way that the wave dash and the tilde are mapped to and from Unicode
KO16KSC5601	KSC5601 16-bit Korean
KO16MSWIN949	Microsoft Windows Code Page 949 Korean
TH8TISASCII	Thai Industrial Standard 620-2533 - ASCII 8-bit
VN8MSWIN1258	Microsoft Windows Code Page 1258 8-bit Vietnamese
ZHS16CGB231280	CGB2312-80 16-bit Simplified Chinese
ZHS16GBK	GBK 16-bit Simplified Chinese
ZHS32GB18030	GB18030-2000
ZHT16BIG5	BIG5 16-bit Traditional Chinese
ZHT16HKSCS	Microsoft Windows Code Page 950 with Hong Kong Supplementary Character Set HKSCS-2001 (character set conversion to and from Unicode is based on Unicode 3.0)
ZHT16MSWIN950	Microsoft Windows Code Page 950 Traditional Chinese
ZHT32EUC	EUC 32-bit Traditional Chinese

European Character Sets

Name	Description
BLT8CP921	Latvian Standard LVS8-92(1) Windows/UNIX 8-bit Baltic
BLT8ISO8859P13	ISO 8859-13 Baltic

Name	Description
BLT8MSWIN1257	Microsoft Windows Code Page 1257 8-bit Baltic
BLT8PC775	IBM-PC Code Page 775 8-bit Baltic
CEL8ISO8859P14	ISO 8859-13 Celtic
CL8ISO8859P5	ISO 8859-5 Latin/Cyrillic
CL8KOI8R	RELCOM Internet Standard 8-bit Latin/Cyrillic
CL8KOI8U	KOI8 Ukrainian Cyrillic
CL8MSWIN1251	Microsoft Windows Code Page 1251 8-bit Latin/Cyrillic
EE8ISO8859P2	ISO 8859-2 East European
EL8ISO8859P7	ISO 8859-7 Latin/Greek
ET8MSWIN923	Microsoft Windows Code Page 923 8-bit Estonian
EE8MSWIN1250	Microsoft Windows Code Page 1250 8-bit East European
EL8MSWIN1253	Microsoft Windows Code Page 1253 8-bit Latin/Greek
EL8PC737	IBM-PC Code Page 737 8-bit Greek/Latin
EE8PC852	IBM-PC Code Page 852 8-bit East European
LT8MSWIN921	Microsoft Windows Code Page 921 8-bit Lithuanian
NE8ISO8859P10	ISO 8859-10 North European
NEE8ISO8859P4	ISO 8859-4 North and North-East European
RU8PC866	IBM-PC Code Page 866 8-bit Latin/Cyrillic
SE8ISO8859P3	ISO 8859-3 South European
US7ASCII	ASCII 7-bit American
US8PC437	IBM-PC Code Page 437 8-bit American
WE8ISO8859P1	ISO 8859-1 West European
WE8ISO8859P15	ISO 8859-15 West European
WE8MSWIN1252	Microsoft Windows Code Page 1252 8-bit West European
WE8PC850	IBM-PC Code Page 850 8-bit West European
WE8PC858	IBM-PC Code Page 858 8-bit West European

Middle Eastern Character Sets

Name	Description
AR8ADOS720	Arabic MS-DOS 720 Server 8-bit Latin/Arabic
AR8ASMO8X	ASMO Extended 708 8-bit Latin/Arabic
AR8ISO8859P6	ISO 8859-6 Latin/Arabic
AR8MSWIN1256	Microsoft Windows Code Page 1256 8-Bit Latin/Arabic
AZ8ISO8859P9E	ISO 8859-9 Latin Azerbaijani
IW8ISO8859P8	ISO 8859-8 Latin/Hebrew
IW8MSWIN1255	Microsoft Windows Code Page 1255 8-bit Latin/Hebrew
TR8MSWIN1254	Microsoft Windows Code Page 1254 8-bit Turkish
TR8PC857	IBM-PC Code Page 857 8-bit Turkish
WE8ISO8859P9	ISO 8859-9 West European & Turkish

TimesTen Character Set

Name	Description
TIMESTEN8	TimesTen legacy character semantics

Universal Character Sets

Name	Descriptions
AL16UTF16	Unicode 4.0 UTF-16 Universal character set. This is the implicit TimesTen national character set. Note: AL16UTF16 is not allowed as a database or connection character set.
AL32UTF8	Unicode 4.0 UTF-8 Universal character set
UTF8	Unicode 3.0 UTF-8 Universal character set, CESU-8 compliant

Supported Linguistic Sorts

The tables in this section list the supported values for the **NLS_SORT** general connection attribute and the NLS_SORT SQL function.

Monolingual Linguistic Sorts

Basic Name	Extended Name
ARABIC	-
ARABIC_MATCH	-
ARABIC_ABJ_SORT	-
ARABIC_ABJ_MATCH	-
ASCII7	-
AZERBAIJANI	XAZERBAIJANI
BENGALI	-
BIG5	-
BINARY	-
BULGARIAN	-
CANADIAN FRENCH	-
CATALAN	XCATALAN
CROATIAN	XCROATIAN
CZECH	XCZECH
CZECH_PUNCTUATION	XCZECH_PUNCTUATION
DANISH	XDANISH
DUTCH	XDUTCH
EBCDIC	-
EEC_EURO	-
EEC_EUROPA3	-
ESTONIAN	-

Basic Name	Extended Name
FINNISH	-
FRENCH	XFRENCH
GERMAN	XGERMAN
GERMAN_DIN	XGERMAN_DIN
GBK	-
GREEK	-
HEBREW	-
HKSCS	-
HUNGARIAN	XHUNGARIAN
ICELANDIC	-
INDONESIAN	-
ITALIAN	-
LATIN	-
LATVIAN	-
LITHUANIAN	-
MALAY	-
NORWEGIAN	-
POLISH	-
PUNCTUATION	XPUNCTUATION
ROMANIAN	-
RUSSIAN	-
SLOVAK	XSLOVAK
SLOVENIAN	XSLOVENIAN
SPANISH	XSPANISH
SWEDISH	-
SWISS	XSWISS

Basic Name	Extended Name
THAI_DICTIONARY	-
TURKISH	XTURKISH
UKRAINIAN	-
UNICODE_BINARY	-
VIETNAMESE	-
WEST_EUROPEAN	XWEST_EUROPEAN

Multilingual Linguistic Sorts

Sort Name	Description
CANADIAN_M	Canadian French sort supports reverse secondary, special expanding characters.
DANISH_M	Danish sort supports sorting uppercase characters before lowercase characters.
FRENCH_M	French sort supports reverse sort for secondary.
GENERIC_M	Generic sorting order which is based on ISO14651 and Unicode canonical equivalence rules but excluding compatible equivalence rules.
JAPANESE_M	Japanese sort supports SJIS character set order and EUC characters which are not included in SJIS.
KOREAN_M	Korean sort Hangul characters are based on Unicode binary order. Hanja characters based on pronunciation order. All Hangul characters are before Hanja characters.
SPANISH_M	Traditional Spanish sort supports special contracting characters.
THAI_M	Thai sort supports swap characters for some vowels and consonants.
SCHINESE_RADICAL_M	Simplified Chinese sort is based on radical as primary order and number of strokes order as secondary order.
SCHINESE_STROKE_M	Simplified Chinese sort uses number of strokes as primary order and radical as secondary order.
SCHINESE_PINYIN_M	Simplified Chinese Pinyin sorting order.

Sort Name	Description
TCHINESE_RADICAL_M	Traditional Chinese sort based on radical as primary order and number of strokes order as secondary order.
TCHINESE_STROKE_M	Traditional Chinese sort uses number of strokes as primary order and radical as secondary order. It supports supplementary characters.

Using the ttIsql Utility

The TimesTen **ttIsql** utility is a general tool for working with a TimesTen data source. The **ttIsql** command line interface is used to execute SQL statements and built-in **ttIsql** commands to perform various operations. Some common tasks that are typically accomplished using **ttIsql** include:

- Data store setup and maintenance. Creating tables and indexes, altering existing tables and updating table statistics can be performed quickly and easily using **ttIsql**.
- Retrieval of information on data store structures. The definitions for tables, indexes and cache groups can be retrieved using built-in **ttIsql** commands. In addition, the current size and state of the data store can be displayed.
- Optimizing data store operations. The **ttIsql** utility can be used to alter and display query optimizer plans for the purpose of tuning SQL operations. The time required to execute various ODBC function calls can also be displayed.

This chapter describes how the **ttIsql** utility is used to perform these types of tasks. The topics are:

- [Batch mode vs. interactive mode](#)
- [Customizing the ttIsql command prompt](#)
- [Using ttIsql's online help](#)
- [Using ttIsql's 'editline' feature \(UNIX only\)](#)
- [Using ttIsql's command history](#)
- [Working with character sets](#)
- [Working with transactions](#)
- [Displaying data store information](#)
- [Viewing and setting data store attributes](#)
- [Viewing and changing query optimizer plans](#)
- [Timing ODBC function calls](#)
- [Working with prepared and parameterized SQL statements](#)
- [Defining default settings with the TTISQL environment variable](#)
- [Managing XLA bookmarks](#)

For more information on TimesTen SQL and for a detailed description of all **ttlsq** commands see the *Oracle TimesTen In-Memory Database API Reference Guide*.

Batch mode vs. interactive mode

The **ttIsql** utility can be used in two distinctly different ways: batch mode or interactive mode. When **ttIsql** is used in interactive mode, users type commands directly into **ttIsql** from the console. When **ttIsql** is used in batch mode, a prepared script of **ttIsql** commands is executed by specifying the name of the file containing the commands.

Batch mode is commonly used for the following types of tasks:

- Performing periodic maintenance operations including the updating of table statistics, compacting the data store and purging log files.
- Initializing a data store by creating tables, indexes and cache groups and then populating the tables with data.
- Generating simple reports by executing common queries.

Interactive mode is suited for the following types of tasks:

- Experimenting with TimesTen features, testing design alternatives and improving query performance.
- Solving data store problems by examining data store statistics.
- Any other data store tasks that are not performed routinely.

By default, when starting **ttIsql** from the shell, **ttIsql** is in interactive mode. The **ttIsql** utility prompts you to type in a valid **ttIsql** built-in command or SQL statement by printing the **Command>** prompt.

Example 5.1

```
C:\>ttIsql
```

```
ttIsql (c) 1996-2006, Oracle. All rights reserved.  
Type ? or "help" for help, type "exit" to quit ttIsql.  
All commands must end with a semicolon character.
```

```
Command>
```

All built-in **ttIsql** commands and SQL statements should be terminated with a semicolon (;) character. The semicolon character tells **ttIsql** that the preceding command is ready to be processed.

Batch mode can be accessed in two different ways. The most common way is to specify the **-f** option on the **ttIsql** command line followed by the name of file to run.

Example 5.2

For example, executing a file containing a CREATE TABLE statement will look like this:

```
C:\>ttIsql -f create.sql MY_DSN
```

```
ttIsql (c) 1996-2006, Oracle. All rights reserved.
```

Type ? or "help" for help, type "exit" to quit ttIsql.
All commands must end with a semicolon character.

```
Command> connect "DSN=MY_DSN;Overwrite=1"
Connection successful: DSN=MY_DSN;DataStore=E:\ds\MY_DSN;
DRIVER=E:\WINNT\System32\TTdv70.dll;
(Default setting AutoCommit=1)
```

```
Command> run "create.sql"
```

```
CREATE TABLE LOOKUP (KEY NUMBER NOT NULL PRIMARY KEY, VALUE CHAR
(64))
```

```
Command> exit
Disconnecting...
Done.
```

```
C:\>
```

The other way to use batch mode is to enter the `run` command directly from the interactive command prompt. The `run` command is followed by the name of the file containing **ttIsql** built-in commands and SQL statements to execute.

Example 5.3 Command> run "create.sql";

```
CREATE TABLE LOOKUP (KEY NUMBER NOT NULL PRIMARY KEY, VALUE CHAR
(64))
Command>
```

Customizing the ttIsql command prompt

You can customize the **ttIsql** command prompt by using the `set` command with the `prompt` attribute.

Example 5.4 `Command> set prompt MY_DSN;`
`MY_DSN`

You can specify a string format (`%c`) that returns the name of the current connection.

Example 5.5 `Command> set prompt %c;`
`con1`

If you want to embed spaces, you must quote the string.

Example 5.6 `Command> set prompt "MY_DSN %c> ";`
`MY_DSN con1>`

Using ttIsql's online help

The **ttIsql** utility has an online version of command syntax definitions and descriptions for all built-in **ttIsql** commands. To access this online help from within **ttIsql** use the **help** command. To view a detailed description of any built-in **ttIsql** commands type the **help** command followed by one or more **ttIsql** commands to display help for. The example below displays the online description for the **connect** and **disconnect** commands.

Example 5.7 Command> help connect disconnect;

All commands must end with a semicolon character.
Arguments in <> are required.
Arguments in [] are optional.

Command Usage: connect [DSN|connection_string]
Command Aliases: (none)
Description: Connects to the data source specified by the optional DSN or connection string argument. If an argument is not given, then the DSN or connection string from the last successful connection is used.
Requires an active connection: NO
Requires autocommit turned off: NO
Reports elapsed execution time: YES
Works only with a TimesTen data source: NO
Example: connect; -or- connect RunData_tt70_32; -or- connect "DSN=RunData_tt70_32;Overwrite=1";

Command Usage: disconnect
Command Aliases: (none)
Description: Disconnects from the currently connected data source. If a transaction is active when disconnecting then the transaction will be rolled back automatically. If a connection exists when executing the "bye", "quit" or "exit" commands then the "disconnect" command will be executed automatically.
Requires an active connection: YES
Requires autocommit turned off: NO
Reports elapsed execution time: YES
Works only with a TimesTen data source: NO
Example: disconnect;
Command>

To view a short description of all **ttIsql** built-in commands type the `help` command without an argument. To view a detailed description of all built-in **ttIsql** commands type the `help` command followed by the `all` argument.

Using ttlsql's 'editline' feature (UNIX only)

On UNIX systems, you can use the 'editline' library to set up emacs (default) or vi bindings that enable you to scroll through previous **ttlsql** commands, as well as edit and resubmit them. This feature is not available or needed on Windows.

To disable the 'editline' feature in **ttlsql**, use the **ttlsql** command `set editline off`.

The set up and keystroke information is described for each type of editor:

- [Emacs binding](#)
- [vi binding](#)

Emacs binding

To get the current settings, create a file `~/.editrc` and put "bind" on the last line of the file, run **ttlsql**. The editline lib will print the current bindings.

The keystrokes when using **ttlsql** with the emacs binding are:

Keystroke	Action
<Left-Arrow>	Move the insertion point left (back up)
<Right-Arrow>	Move the insertion point right (forward)
<Up-Arrow>	Scroll to the command prior to the one being displayed. Places the cursor at the end of the line.
<Down-Arrow>	Scroll to a more recent command history item and put the cursor at the end of the line.
<Ctrl-A>	Move the insertion point to the beginning of the line.
<Ctrl-E>	Move the insertion point to the end of the line.
<Ctrl-K>	"Kill" (Save and erase) the characters on the command line from the current position to the end of the line.
<Ctrl-Y>	"Yank" (Restore) the characters previously saved and insert them at the current insertion point.
<Ctrl-F>	Forward char - move forward 1 (see Right Arrow)
<Ctrl-B>	Backward char - move back 1 (see Left Arrow)
<Ctrl-P>	Previous History (see Up Arrow)
<Ctrl-N>	Next History (see up Down Arrow)

vi binding

To use the vi bindings, create a file `~/.editrc` and put “bind-v” in the file, run **ttIsql**. To get the current settings, create a file `~/.editrc` and put “bind” on the last line of the file. When you execute **ttIsql**, the editline lib will print the current bindings.

The keystrokes when using **ttIsql** with the vi binding are:

Keystroke	Action
<Left-Arrow>, h	Move the insertion point left (back up)
<Right-Arrow>, l	Move the insertion point right (forward)
<Up-Arrow>, k	Scroll to the prior command in the history and put the cursor at the end of the line.
<Down-Arrow>, j	Scroll to the next command in the history and put the cursor at the end of the line.
ESC	Vi Command mode
0, \$	Move the insertion point to the beginning of the line, Move to end of the line.
i, I	Insert mode, Insert mode at beginning of the line
a, A	Add (“Insert after”) mode, Append at end of line
R	Replace mode
C	Change to end of line
B	Move to previous word
e	Move to end of word
<Ctrl-P>	Previous History (see Up Arrow)
<Ctrl-N>	Next History (see up Down Arrow)

Using ttlsql's command history

The **ttlsql** utility stores a list of the last 100 commands executed within the current **ttlsql** session. The commands in this list can be viewed or executed again without having to type the entire command over. Both SQL statements and built-in **ttlsql** commands are stored in the history list. Use the `history` command (or `h` for short) to view the list of previously executed commands.

Example 5.8

```
Command> h;
8      INSERT INTO T3 VALUES (3)
9      INSERT INTO T1 VALUES (4)
10     INSERT INTO T2 VALUES (5)
11     INSERT INTO T3 VALUES (6)
12     autocommit 0
13     showplan
14     SELECT * FROM T1, t2, t3 WHERE A=B AND B=C AND A=B
15     trytbllocks 0
16     tryserial 0
17     SELECT * FROM T1, t2, t3 WHERE A=B AND B=C AND A=B
Command>
```

The `history` command displays the last 10 SQL statements or **ttlsql** built-in commands executed. To display more than that last 10 commands specify the maximum number to display as an argument to the `history` command.

Each entry in the history list is identified by a unique number. The `!` character followed by the number of the command can be used to execute the command again.

For example:

Example 5.9

```
Command>
Command> ! 12;
```

```
autocommit 0
Command>
```

To execute the last command again simply type a sequence of two `!` characters.

```
Command> !!;

autocommit 0
Command>
```

To execute the last command that begins with a given string type the **!** character followed by the first few letters of the command. For example:

Example 5.10 Command> ! auto;

```
autocommit 0
Command>
```

Saving and clearing `ttIsql`'s command history

You can save the list of commands that `ttIsql` stores by using the `savehistory` command.

Example 5.11 Command> savehistory history.txt;

If the output file already exists, use the `-a` option to append the new command history to the file or the `-f` option to overwrite the file. The next example shows how to append new command history to an existing file.

Example 5.12 Command> savehistory -a history.txt;

You can clear the list of commands that `ttIsql` stores by using the `clearhistory` command.

Example 5.13 Command> clearhistory;

Working with character sets

The **ttIsql** utility supports the character sets listed in [Chapter 4, “Supported Character Sets.”](#) The ability of **ttIsql** to display characters depends on the native OS locale settings of the terminal on which you are using **ttIsql**.

To override the locale-based output format, use the `ncharencoding` option or the `-N` option. The valid values for these options are `LOCALE` (the default) and `ASCII`. If you choose `ASCII` and **ttIsql** encounters a Unicode character, it displays it in escaped format.

You do not need to have an active connection to change the output method.

Working with transactions

The **ttIsql** utility has several built-in commands for managing transactions. These commands are summarized below:

- **autocommit** – Turns on or off the autocommit feature.
- **commit** – Commits the current transaction.
- **commitdurable** – Commits the current transaction and ensures that the committed work will be recovered in case of data store failure.
- **rollback** – Rolls back the current transaction.
- **isolation** – Changes the transaction isolation level.
- **sqlquerytimeout** – Specifies the number of seconds to wait for a SQL statement to execute before returning to the application.

When starting **ttIsql** the autocommit feature is turned on by default. In this mode every SQL operation against the data store is committed automatically. To turn the autocommit feature off execute **ttIsql**'s built-in **autocommit** command with an argument of 0.

When autocommit is turned off transactions must be committed or rolled back manually by executing **ttIsql**'s **commit**, **commitdurable** or **rollback** commands. The **commitdurable** command will ensure that the transaction's effect is preserved in case of data store failure.

The **isolation** command can be used to change the current connection's transaction isolation properties. The isolation can be changed only at the beginning of a transaction. The **isolation** command accepts one of the following constants: **READ_COMMITTED** and **SERIALIZABLE**. If the **isolation** command is executed without an argument then the current isolation level is reported.

The **sqlquerytimeout** command sets the timeout period for SQL statements. If the execution time of a SQL statement exceeds the number of seconds set by **sqlquerytimeout**, the SQL statement is not executed and an 6111 error is generated. For details, see "[Setting a timeout value for executing SQL statements](#)" in the *Oracle TimesTen In-Memory Database Java Developer's and Reference Guide* and "[Setting a timeout value for executing SQL statements](#)" in the *Oracle TimesTen In-Memory Database C Developer's and Reference Guide*.

Note: TimesTen rollback and query timeout features do not stop Cache Connect operations that are being processed on Oracle. This includes passthrough statements, flushing, manual loading, manual refreshing, synchronous writethrough, propagating, and transparent loading.

The following example demonstrates the common use of **ttIsql**'s built-in transaction management commands.

Example 5.14

```
E:\>ttIsql
ttIsql (c) 1996-2006, Oracle. All rights reserved.
Type ? or "help" for help, type "exit" to quit ttIsql.
All commands must end with a semicolon character.

Command> connect "DSN=MY_DSN";
Connection successful:
DSN=MY_DSN;DataStore=E:\ds\MY_DSN;DRIVER=E:\WINNT\System32\
TTdv70.dll;
(Default setting AutoCommit=1)
Command> autocommit 0;
Command> CREATE TABLE LOOKUP (KEY NUMBER NOT NULL PRIMARY KEY,
VALUE CHAR (64));
Command> commit;
Command> INSERT INTO LOOKUP VALUES (1, 'ABC');
1 row inserted.
Command> SELECT * FROM LOOKUP;
< 1,
ABC
>
1 row found.
Command> rollback;
Command> SELECT * FROM LOOKUP;
0 rows found.
Command> isolation;
isolation = READ_COMMITTED
Command> commitdurable;
Command> sqlquerytimeout 10;
Command> sqlquerytimeout;
Query timeout = 10 seconds
Command> disconnect;
Disconnecting...
Command> exit;
Done.
E:\>
```

Displaying data store information

There are several built-in **ttIsql** commands that display information on data store structures. The most useful commands are summarized below:

- **describe** – Displays information on tables, prepared statements and procedures.
- **cachegroups** – Displays the attributes of cache groups.
- **dssize** – Reports the current sizes of the permanent and temporary data store partitions.
- **monitor** – Displays a summary of the current state of the data store.

The `describe` command is used to display information on table and result set columns as well as parameters for prepared SQL statements and built-in procedures. The argument to the `describe` command can be the name of a table, a built-in procedure, a SQL statement or a command id for a previously prepared SQL statement.

Example 5.15

```
Command> CREATE TABLE T1 (KEY NUMBER NOT NULL PRIMARY KEY, VALUE
CHAR (64));
Command> describe T1;
```

```
Table USER.T1:
```

```
Columns:
```

*KEY	NUMBER NOT NULL
VALUE	CHAR (64)

```
1 table found.
```

```
(primary key columns are indicated with *)
```

```
Command> describe SELECT * FROM T1 WHERE KEY=?;
```

```
Prepared Statement:
```

```
Parameters:
```

Parameter 1	NUMBER
-------------	--------

```
Columns:
```

KEY	NUMBER NOT NULL
VALUE	CHAR (64)

```
Command> describe ttOptUseIndex;
```

```
Procedure TTOPTUSEINDEX:
```

```
Parameters:
```

Parameter INDOPTION	VARCHAR (1024)
---------------------	----------------

```
Columns:
```

```
(none)
```

```
1 procedure found.
```

```
Command>
```

The `cachegroups` command is used to provide detailed information on cache groups defined in the current data store. The attributes of the root and child tables defined in the cache group are displayed in addition to the WHERE clauses associated with the cache group and the DURATION value for cache groups that use cache aging. The argument to the `cachegroups` command is the name of the cache group that you want to display information for.

Example 5.16 `Command> cachegroups MY_CACHE_GROUP;`

Cache Group USER.MY_CACHE_GROUP:

Duration: 40 Minutes

Root Table: USER.T1

Where Clause: (T1.KEY < 100)

Type: Not Propagate

Child Table: USER.T2

Where Clause: (none)

Type: Propagate

1 cache group found.

`Command>`

The `dssize` command is used to report the current memory status of the permanent and temporary partitions as well as the maximum, allocated and in-use sizes for the data store. The `monitor` command displays all of the information provided by the `dssize` command plus additional statistics on the number of connections, checkpoints, lock timeouts, commits, rollbacks and other information collected since the last time the data store was loaded into memory.

Example 5.17 `Command> monitor;`

TIME_OF_1ST_CONNECT:	Thu Sep 23 11:51:38 2005
DS_CONNECTS:	4
DS_DISCONNECTS:	0
DS_CHECKPOINTS:	2
DS_CHECKPOINTS_FUZZY:	0
DS_COMPACTS:	0
PERM_ALLOCATED_SIZE:	204800
PERM_IN_USE_SIZE:	780
PERM_IN_USE_HIGH_WATER:	780
TEMP_ALLOCATED_SIZE:	36864
TEMP_IN_USE_SIZE:	2048
TEMP_IN_USE_HIGH_WATER:	2048
SYS18:	0
XACT_BEGINS:	16
XACT_COMMITS:	15


```
XACT_D_COMMITS:          0
XACT_ROLLBACKS:         0
LOG_FORCES:              2
DEADLOCKS:               0
LOCK_TIMEOUTS:           0
LOCK_GRANTS_IMMED:      290
LOCK_GRANTS_WAIT:        0
CMD_PREPARES:            15
CMD_REPREPARES:          0
CMD_TEMP_INDEXES:        0
LAST_LOG_FILE:           0
REPHOLD_LOG_FILE:        -1
REPHOLD_LOG_OFF:         -1
REP_XACT_COUNT:           0
REP_CONFLICT_COUNT:      0
REP_PEER_CONNECTIONS:    0
REP_PEER_RETRIES:        0
FIRST_LOG_FILE:          0
LOG_BYTES_TO_LOG_BUFFER: 10960
LOG_FS_READS:             0
LOG_FS_WRITES:           2
LOG_BUFFER_WAITS:         0
CHECKPOINT_BYTES_WRITTEN: 1037176
SYS1:                     0
SYS2:                     0
SYS3:                     0
SYS4:                     0
SYS5:                     0
SYS6:                     0
SYS7:                     23
SYS8:                     0
SYS10:                    0
SYS11:                    0
SYS12:                    0
SYS13:                    0
SYS14:                    0
SYS15:                    0
SYS16:                    0
SYS17:                    0
SYS19:                    0
SYS9:
```

Command>

Viewing and setting data store attributes

You can view and set data store attributes with the **ttIsql** `show` and `set` commands. For a list of the attributes that you can view and set with **ttIsql**, see "Set/Show attributes" in *Oracle TimesTen In-Memory Database API Reference Guide*.

Example 5.18 To view the setting for the **PassThrough** attribute, enter the following command:

```
Command> show passthrough;  
PassThrough = 0
```

To change the **PassThrough** setting, **AutoCommit** must be 0. Change the **PassThrough** setting as follows:

Example 5.19

```
Command> autocommit=0;  
Command> set passthrough 1;
```

Viewing and changing query optimizer plans

The built-in `showplan` command is used to display the query optimizer plans used by the TimesTen Data Manager for executing queries. In addition, **ttIsql** contains built-in query optimizer hint commands for altering the query optimizer plan. By using the `showplan` command in conjunction with the built-in commands summarized below, the optimum execution plan can be designed. For detailed information on the TimesTen query optimizer see [Chapter 9](#) the *Oracle TimesTen In-Memory Database C Developer's and Reference Guide*.

- **optprofile** – Displays the current optimizer hint settings and join order.
- **setjoinorder** – Sets the join order.
- **setuseindex** – Sets the index hint.
- **tryhash** – Enables or disables the use of hash indexes.
- **trymergejoin** – Enables or disables merge joins.
- **trynestedloopjoin** – Enables or disables nested loop joins.
- **tryserial** – Enables or disables serial scans.
- **trytmphash** – Enables or disables the use of temporary hash indexes.
- **trytmptable** – Enables or disables the use of an intermediate results table.
- **trytmpttree** – Enables or disables the use of temporary ttree indexes.
- **tryttree** – Enables or disables the use of ttree indexes.
- **tryrowid** – Enables or disables the use of rowid scans.
- **trytbllocks** – Enables or disables the use of table locks.
- **unsetjoinorder** – Clears the join order.
- **unsetuseindex** – Clears the index hint.

When using the `showplan` command and the query optimizer hint commands the autocommit feature must be turned off. Use **ttIsql**'s `autocommit` built-in command to turn autocommit off.

The example below shows how these commands can be used to change the query optimizer execution plan.

Example 5.20

```
Command> CREATE TABLE T1 (A NUMBER);
Command> CREATE TABLE T2 (B NUMBER);
Command> CREATE TABLE T3 (C NUMBER);
Command>
Command> INSERT INTO T1 VALUES (3);
1 row inserted.
Command> INSERT INTO T2 VALUES (3);
1 row inserted.
Command> INSERT INTO T3 VALUES (3);
1 row inserted.
Command> INSERT INTO T1 VALUES (4);
1 row inserted.
```

```

Command> INSERT INTO T2 VALUES (5);
1 row inserted.
Command> INSERT INTO T3 VALUES (6);
1 row inserted.
Command>
Command> autocommit 0;
Command> showplan;
Command> SELECT * FROM T1, T2, T3 WHERE A=B AND B=C AND A=B;

```

Query Optimizer Plan:

```

STEP:          1
LEVEL:         3
OPERATION:     TblLkSerialScan
TBLNAME:       T1
IXNAME:        <NULL>
PRED:          <NULL>
OTHERPRED:     <NULL>

STEP:          2
LEVEL:         3
OPERATION:     TblLkSerialScan
TBLNAME:       T2
IXNAME:        <NULL>
PRED:          <NULL>
OTHERPRED:     T1.A = T2.B AND T1.A = T2.B

STEP:          3
LEVEL:         2
OPERATION:     NestedLoop
TBLNAME:       <NULL>
IXNAME:        <NULL>
PRED:          <NULL>
OTHERPRED:     <NULL>

STEP:          4
LEVEL:         2
OPERATION:     TblLkSerialScan
TBLNAME:       T3
IXNAME:        <NULL>
PRED:          <NULL>
OTHERPRED:     T2.B = T3.C

STEP:          5
LEVEL:         1
OPERATION:     NestedLoop
TBLNAME:       <NULL>
IXNAME:        <NULL>
PRED:          <NULL>

```

```
OTHERPRED: <NULL>

< 3, 3, 3 >
1 row found.
Command> trytbllocks 0;
Command> tryserial 0;
Command> SELECT * FROM T1, t2, t3 WHERE A=B AND B=C AND A=B;
```

Query Optimizer Plan:

```
STEP:          1
LEVEL:         3
OPERATION:     TmpTtreeScan
TBLNAME:      T1
IXNAME:       <NULL>
PRED:         <NULL>
OTHERPRED:    <NULL>

STEP:          2
LEVEL:         3
OPERATION:     TmpTtreeScan
TBLNAME:      T2
IXNAME:       <NULL>
PRED:         T2.B >= T1.A
OTHERPRED:    <NULL>

STEP:          3
LEVEL:         2
OPERATION:     MergeJoin
TBLNAME:      <NULL>
IXNAME:       <NULL>
PRED:         T1.A = T2.B AND T1.A = T2.B
OTHERPRED:    <NULL>

STEP:          4
LEVEL:         2
OPERATION:     TmpTtreeScan
TBLNAME:      T3
IXNAME:       <NULL>
PRED:         <NULL>
OTHERPRED:    T2.B = T3.C

STEP:          5
LEVEL:         1
OPERATION:     NestedLoop
TBLNAME:      <NULL>
IXNAME:       <NULL>
PRED:         <NULL>
```

```
OTHERPRED: <NULL>

< 3, 3, 3 >
1 row found.
Command>
```

In this example a query against three tables is executed and the query optimizer plan is displayed. The first version of the query simply uses the query optimizer's default execution plan. However, in the second version the `trytbllocks` and `tryserial` built-in hint commands have been used to alter the query optimizer's plan. Instead of using serial scans and nested loop joins the second version of the query uses temporary index scans and merge joins.

In this way the `showplan` command in conjunction with `ttIsql`'s built-in query optimizer hint commands can be used to quickly determine which execution plan should be used to meet application requirements.

Timing ODBC function calls

Information on the time required to execute common ODBC function calls can be displayed by using **ttIsql**'s `timing` command. When the timing feature is enabled many built-in **ttIsql** commands will report the elapsed execution time associated with the primary ODBC function call corresponding to the **ttIsql** command that is executed.

For example, when executing **ttIsql**'s `connect` command several ODBC function calls are executed, however, the primary ODBC function call associated with `connect` is **SQLDriverConnect** and this is the function call that is timed and reported as shown below.

Example 5.21

```
Command> timing 1;
Command> connect "DSN=MY_DSN";
Connection successful:
DSN=MY_DSN;DataStore=E:\ds\MY_DSN;DRIVER=E:\WINNT\System32\
    TTdv70.dll;
(Default setting AutoCommit=1)
Execution time (SQLDriverConnect) = 1.2626 seconds.
Command>
```

In the example above, the **SQLDriverConnect** call took about 1.25 seconds to execute.

When using the timing command to measure queries, the time required to execute the query plus the time required to fetch the query results is measured. To avoid measuring the time to format and print query results to the display, set the verbosity level to 0 before executing the query.

Example 5.22

```
Command> timing 1;
Command> verbosity 0;
Command> SELECT * FROM T1;
Execution time (SQLExecute + FetchLoop) = 0.064210 seconds.
Command>
```

Working with prepared and parameterized SQL statements

Preparing a SQL statement just once and then executing it multiple times is much more efficient for TimesTen applications than re-preparing the statement each time it is to be executed. **ttIsql** has a set of built-in commands to work with prepared SQL statements. These commands are summarized below:

- **prepare** – Prepares a SQL statement. Corresponds to a **SQLPrepare** ODBC call.
- **exec** – Executes a previously prepared statement. Corresponds to a **SQLExecute** ODBC call.
- **execandfetch** – Executes a previously prepared statement and fetches all result rows. Corresponds to a **SQLExecute** call followed by one or more calls to **SQLFetch**.
- **fetchall** – Fetches all result rows for a previously executed statement. Corresponds to one or more **SQLFetch** calls.
- **fetchone** – Fetches only one row for a previously executed statement. Corresponds to exactly one **SQLFetch** call.
- **close** – Closes the result set cursor on a previously executed statement that generated a result set. Corresponds to a **SQLFreeStmt** call with the **SQL_CLOSE** option.
- **free** – Closes a previously prepared statement. Corresponds to a **SQLFreeStmt** call with the **SQL_DROP** option.
- **describe** – Describes the prepared statement including the input parameters and the result columns.

The **ttIsql** utility prepared statement commands also handle SQL statement parameter markers. When parameter markers are included in a prepared SQL statement, **ttIsql** will automatically prompt for the value of each parameter in the statement at execution time.

The example below uses **ttIsql**'s prepared statement commands to prepare an **INSERT** statement into a table containing an **NUMBER** and a **CHAR** column. The statement is prepared and then executed twice with different values for each of the statement's two parameters. **ttIsql**'s **timing** command is used to display the elapsed time required to executed the primary ODBC function call associated with each command.

Example 5.23

```
Command> connect "DSN=MY_DSN;Overwrite=1";
Connection successful:
DSN=MY_DSN;DataStore=E:\ds\MY_DSN;Overwrite=1;DRIVER=E:\WINNT\System32\Ttdv70.dll;
(Default setting AutoCommit=1)
Command> timing 1;
Command> CREATE TABLE T1 (KEY NUMBER NOT NULL PRIMARY KEY, VALUE
CHAR (64));
```



```
Execution time (SQLExecDirect) = 0.1337 seconds.  
Command> prepare INSERT INTO T1 VALUES (?,?);  
Execution time (SQLPrepare) = 0.0668 seconds.  
Command> exec;
```

All parameter values must be terminated with a semicolon character.

Type '?' for help on entering parameter values.

Type '*' to abort the parameter entry process.

```
Enter Parameter 1 (NUMBER) >1;  
Enter Parameter 2 (CHAR) >'abc';  
1 row inserted.  
Execution time (SQLExecute) = 0.0757 seconds.  
Command> exec;
```

All parameter values must be terminated with a semicolon character.

Type '?' for help on entering parameter values.

Type '*' to abort the parameter entry process.

```
Enter Parameter 1 (NUMBER) >2;  
Enter Parameter 2 (CHAR) >'def';  
1 row inserted.  
Execution time (SQLExecute) = 0.0306 seconds.  
Command> free;  
Command> SELECT * FROM T1;  
< 1,  
abc  
>  
< 2,  
def  
>  
2 rows found.  
Execution time (SQLExecDirect) = 0.0316 seconds.  
Command> disconnect;  
Disconnecting...  
Execution time (SQLDisconnect) = 1.7091 seconds.  
Command>  
Command>
```

In the example above, the `prepare` command is immediately followed by the SQL statement to prepare. Whenever a SQL statement is prepared in **ttIsql** a unique command ID is assigned to the prepared statement. **ttIsql** uses this ID to keep track of multiple prepared statements. A maximum of 256 prepared statements can exist in a **ttIsql** session simultaneously. When the `free` command is executed, the command ID is automatically disassociated from the prepared SQL statement.

To see the command IDs generated by **ttIsql** when using the prepared statement commands, set the verbosity level to 4 using the `verbosity` command before preparing the statement, or use the `describe *` command to list all prepared statements with their IDs.

Command IDs can be referenced explicitly when using **ttIsql**'s prepared statement commands. For a complete description of the syntax of **ttIsql**'s prepared statement commands see the *Oracle TimesTen In-Memory Database API Reference Guide* or type `help` at the **ttIsql** command prompt.

The example below prepares and executes a **SELECT** statement with a predicate containing one **NUMBER** parameter. The `fetchone` command is used to fetch the result row generated by the statement. The `showplan` command is used to display the execution plan used by the TimesTen query optimizer when the statement is executed. In addition, the verbosity level is set to 4 so that the command ID used by **ttIsql** to keep track of the prepared statement is displayed.

Example 5.24

```
Command> connect "DSN=MY_DSN;Overwrite=1";
Connection successful:
DSN=MY_DSN;DataStore=E:\ds\MY_DSN;Overwrite=1;DRIVER=E:\WINNT\System32\TTdv70.dll;
(Default setting AutoCommit=1)
The command succeeded.
Command> CREATE TABLE T1 (KEY NUMBER NOT NULL PRIMARY KEY, VALUE
CHAR (64));
The command succeeded.
Command> INSERT INTO T1 VALUES (1, 'abc');
1 row inserted.
The command succeeded.
Command> autocommit 0;
The command succeeded.
Command> showplan 1;
The command succeeded.
Command> verbosity 4;
The command succeeded.
Command> prepare SELECT * FROM T1 WHERE KEY=?;
Assigning new prepared command id = 0.
```

Query Optimizer Plan:

```
STEP:          1
LEVEL:         1
OPERATION:     RowLkHashScan
TBLNAME:       T1
IXNAME:        T1
PRED:          T1.KEY = qmark_1
OTHERPRED:     <NULL>
```

The command succeeded.

```
Command> exec;  
Executing prepared command id = 0.
```

All parameter values must be terminated with a semicolon character.

Type '?' for help on entering parameter values.

Type '*' to abort the parameter entry process.

```
Enter Parameter 1 (NUMBER) >1;  
The command succeeded.  
Command> fetchone;  
Fetching prepared command id = 0.  
< 1,  
abc  
>  
1 row found.  
The command succeeded.  
Command> close;  
Closing prepared command id = 0.  
The command succeeded.  
Command> free;  
Freeing prepared command id = 0.  
The command succeeded.  
Command> commit;  
The command succeeded.  
Command> disconnect;  
Disconnecting...  
The command succeeded.  
Command>
```

Defining default settings with the TTISQL environment variable

The **ttIsql** utility can be customized to automatically execute a set of command line options every time a **ttIsql** session is started from the command prompt. This is accomplished by setting an environment variable called **TTISQL** to the value of the **ttIsql** command line that you prefer. A summary of **ttIsql** command line options is shown below. For a complete description of **ttIsql**'s command line options see the *Oracle TimesTen In-Memory Database API Reference Guide*.

Example 5.25

```
Usage: ttIsql [-h | -help | -helpcmds | -helpfull | -V]
            [-connStr <connection_string>]
            [-f <filename>]
            [-v <verbosity>]
            [-e <initialization_commands>]
            [-interactive]
            [-N <ncharencoding>]
            [-wait]
```

The **TTISQL** environment variable has the same syntax requirements as the **ttIsql** command line. When **ttIsql** starts up it reads the value of the **TTISQL** environment variable and applies all options specified by the variable to the current **ttIsql** session. If a particular command line option is specified in both the **TTISQL** environment variable and the command line then the command line version will always take precedence.

Example 5.26

The procedure for setting the value of an environment variable differs based on the platform and shell that **ttIsql** is started from. As an example, setting the **TTISQL** environment variable on Windows could look like this:

```
C:\>set TTISQL=-connStr "DSN=MY_DSN" -e "autocommit 0;tables;"
```

Example 5.27

In this example, **ttIsql** will automatically connect to a DSN called **MY_DSN**, turn off autocommit and display all tables in the data store as shown below:

```
C:\>ttIsql

ttIsql (c) 1996-2006, Oracle. All rights reserved.
Type ? or "help" for help, type "exit" to quit ttIsql.
All commands must end with a semicolon character.

Command> connect "DSN=MY_DSN";
Connection successful:
DSN=MY_DSN;DataStore=E:\ds\MY_DSN;DRIVER=E:\WINNT\System32\TTdv70
.dll;
(Default setting AutoCommit=1)
```

```
Command> autocommit 0;

Command> tables;
SYS.CACHE_GROUP
SYS.COLUMNS
SYS.COL_STATS
SYS.INDEXES
SYS.MONITOR
SYS.PLAN
SYS.SEQUENCES
SYS.TABLES
SYS.TBL_STATS
SYS.TRANSACTION_LOG_API
SYS.VIEWS
TTREP.REPELEMENTS
TTREP.REPLICATIONS
TTREP.REPPEERS
TTREP.REPSTORES
TTREP.REPSUBSCRIPTIONS
TTREP.REPTABLES
TTREP.TTSTORES
16 tables found.
Command>
```

Managing XLA bookmarks

You can use the **xldeletebookmark** command to both check the status of the current XLA bookmarks and delete them. This command requires ADMIN privilege or object ownership.

For example, when running the XLA application, 'xlaSimple,' you can check the bookmark status by entering:

```
Command> xldeletebookmark;
```

```
XLA Bookmark: xlaSimple
  Read Log File: 0
  Read Offset:   630000
  Purge Log File: 0
  Purge Offset:  629960
  PID:           2808
  In Use:        No
1 bookmark found.
```

To delete the bookmark, enter:

```
Command> xldeletebookmark xlaSimple;
Command>
```

Working with Data in a TimesTen Data Store

This chapter provides detailed information on the basic components in a TimesTen data store and simple examples of how you can use SQL to manage these components. For more information about SQL, see the *Oracle TimesTen In-Memory Database SQL Reference Guide*.

For information on how to execute SQL from within a C or Java application, see "Managing TimesTen data" in the *Oracle TimesTen In-Memory Database Java Developer's and Reference Guide* or "Managing TimesTen data" in the *Oracle TimesTen In-Memory Database C Developer's and Reference Guide*.

This chapter includes the following topics:

- [Data store overview](#)
- [Understanding tables](#)
- [Working with tables](#)
- [Implementing aging in your tables](#)
- [Understanding materialized views](#)
- [Working with materialized views](#)
- [Understanding views](#)
- [Working with views](#)
- [Understanding indexes](#)
- [Working with indexes](#)
- [Understanding rows](#)
- [Working with rows](#)

Data store overview

This section describes the main TimesTen data store elements and features. It includes the following topics:

- [Data store components](#)
- [Data store users and owners](#)
- [Data store persistence](#)

Data store components

A TimesTen data store has the following permanent components:

- **Tables.** The primary components of a TimesTen data store are the tables that contain the application data. See [“Understanding tables” on page 127](#).
- **Materialized Views.** Read-only tables that hold a summary of data selected from one or more “regular” TimesTen tables. See [“Understanding materialized views” on page 135](#).
- **Views.** Logical tables that are based on one or more tables called *detail tables*. A view itself contains no data. See [“Understanding views” on page 140](#).
- **Indexes.** Indexes on one or more columns of a table may be created for faster access to tables. See [“Understanding indexes” on page 142](#).
- **Rows.** Every table consists of 0 or more rows. A row is a formatted list of values. See [“Understanding rows” on page 144](#).
- **System tables.** System tables contain TimesTen metadata, such as a table of all tables. See [“System tables” on page 129](#) and the chapter ["System and Replication Tables"](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*.

There are also many temporary components, including prepared commands, cursors and locks.

Data store users and owners

Unless Access Control is enabled, the TimesTen Data Manager does not authenticate user names. It accepts user names and ignores passwords entirely. TimesTen Client/Server does authenticate users with passwords. Applications should choose one UID for the application itself because by default the login name that is being used to run the application becomes the owner of the data store. If two different logins are used, TimesTen may have difficulty finding the correct tables. If you omit the UID connection attribute in the connection string, TimesTen uses the current user’s login name. TimesTen converts all user names to upper case characters.

Users cannot access TimesTen data stores as user SYS. TimesTen determines the user name by the value of the UID connection attribute, or if not present, then by the login name of the connected user. If a user’s login is SYS, set the UID connection to override the login name.

Data store persistence

When a data store is created, it has either the permanent or temporary attribute set:

- **Permanent data stores** are stored to disk automatically through a procedure called checkpointing. TimesTen automatically performs background checkpoints based on the settings of the data store attributes [CkptFrequency](#)

and **CkptLogVolume**. TimesTen also checkpoints the data store when the last application disconnects. Applications can also checkpoint a data store directly to disk by invoking the **ttCkpt** or **ttCkptBlocking** built-in procedures described in the *Oracle TimesTen In-Memory Database API Reference Guide*.

- **Temporary data stores** are not stored to disk. A temporary data store is automatically destroyed when no applications are connected to it. TimesTen removes all disk-based files, when the last application disconnects.

Note: You cannot change the permanent or temporary attribute on a data store after it is created.

Understanding tables

A TimesTen table consists of rows that have a common format or structure. This format is described by the table's columns. Each column has:

- A data type
- Optional nullability, primary key and foreign key properties
- An optional default value

Unless you explicitly declare a column NOT NULL, columns are nullable. If a column in a table is nullable, it can contain a NULL value. Otherwise, each row in the table must have a non-NULL value in that column.

The format of TimesTen columns cannot be altered. It is possible to add or remove columns but not to change column definitions. To add or remove columns, use the **ALTER TABLE** statement. To change column definitions, an application must first drop the table and then recreate it with the new definitions.

The rest of this section includes the following topics:

- [In-line and out-of-line columns](#)
- [Default column values](#)
- [Table names](#)
- [Table access](#)
- [Primary keys, foreign keys and unique indexes](#)
- [System tables](#)

In-line and out-of-line columns

The in-memory layout of the rows of a table is designed to provide fast access to rows while minimizing wasted space. TimesTen designates each VARBINARY, NVARCHAR and VARCHAR column of a table as either *in-line* or *not inline*.

- An in-line column has a fixed length. All values of fixed-length columns of a table are stored row wise.

- A not inline column has a varying length. Some VARCHAR, NVARCHAR or VARBINARY data type columns are stored not inline. Not inline columns are not stored contiguously with the row but are allocated. Accessing out-of-line columns is slightly slower than accessing in-line columns. By default, VARCHAR, NVARCHAR and VARBINARY columns whose declared column length is > 128 bytes are stored out of line. Columns whose declared column length is <= 128 bytes are stored inline.

The maximum sizes of in-line and out-of-line portions of a row are listed in [“Estimating table size” on page 130](#).

Default column values

When you create a table, you can specify default values for the columns. The default value you specify must be compatible with the data type of the column. You can specify one of the following default values for a column:

- NULL (for any column type)
- A constant value
- SYSDATE (for DATE and TIMESTAMP columns)
- USER (for CHAR columns)
- CURRENT_USER (for CHAR columns)
- SYSTEM_USER (for CHAR columns)

If you use the DEFAULT clause of the [CREATE TABLE](#) statement but do not specify the default value, the default value is NULL. See ["Column Definition"](#) in [Oracle TimesTen In-Memory Database SQL Reference Guide](#).

Table names

A TimesTen table is identified uniquely by its owner name and table name. Every table has an owner. By default, the owner is the user who created the table. Tables created by TimesTen, such as system tables, have the owner name SYS, or TTREP if created during replication.

To uniquely refer to a table, specify both its owner and name separated by a period “.” (for example, MARY.PAYROLL). If an application does not specify an owner, TimesTen looks for the table under the user name of the caller, then under the user name SYS.

A name is an alphanumeric value that begins with a letter (not a digit). A name can include underscores. The maximum length of a table name is 30 characters. The maximum length of an owner name is also 30 characters. TimesTen displays all table, column and owner names to upper case characters. See [Chapter 2, “Names”](#) in the [Oracle TimesTen In-Memory Database SQL Reference Guide](#) for additional information.

Table access

Applications access tables through SQL statements. The TimesTen query optimizer automatically chooses a fast way to access tables. It uses existing indexes or, if necessary, creates temporary indexes to speed up access. For improved performance, applications should explicitly create indexes for frequently searched columns because the automatic creation and destruction of temporary indexes incurs a performance overhead. For more details, see [“Tune statements and use indexes” on page 174](#).

Primary keys, foreign keys and unique indexes

The creator of a TimesTen table can designate one or more columns as a *primary key* to indicate that duplicate values for that set of columns should be rejected. Primary key columns cannot be nullable. A table can have at most one primary key. TimesTen automatically creates a t-tree index on the primary key to enforce uniqueness on the primary key and to guarantee fast access through the primary key. Indexes are discussed in [“Understanding indexes” on page 142](#). Once a row is inserted, its primary key columns cannot be modified, except to change the t-tree index to a hash index.

Although a table may have only one primary key, additional uniqueness properties may be added to the table using *unique indexes* (see ["CREATE INDEX"](#) in the *Oracle TimesTen In-Memory Database SQL Reference Guide*).

Note: Columns of a primary key cannot be nullable; a unique index can be built on nullable columns.

A table may also have one or more foreign keys through which rows correspond to rows in another table. Foreign keys relate to a primary key or uniquely indexed columns in the other table. Foreign keys use a T-tree index on the referencing columns (see ["CREATE TABLE"](#) in the *Oracle TimesTen In-Memory Database SQL Reference Guide*).

System tables

In addition to tables created by applications, a TimesTen data store contains system tables. System tables contain TimesTen metadata such as descriptions of all tables and indexes in the data store, as well as other information such as optimizer plans. Applications may query system tables just as they query user tables. Applications may not update system tables. TimesTen system tables are described in the chapter ["System and Replication Tables"](#) in the *Oracle TimesTen In-Memory Database SQL Reference Guide*.

Note: TimesTen system table formats may change between releases and are different between the 32- and 64-bit versions of TimesTen.

Working with tables

This section includes the following topics:

- [Creating a table](#)
- [Dropping a table](#)
- [Estimating table size](#)

Creating a table

To create a table, use the SQL statement [CREATE TABLE](#). The syntax for all SQL statements is provided in the [Oracle TimesTen In-Memory Database SQL Reference Guide](#). TimesTen converts table names to upper case characters.

Example 6.1 The following SQL statement creates a table, called *NameID*, with two columns: *CustId* and *CustName* of two different data types.

```
CREATE TABLE NameID (CustId TT_INTEGER, CustName VARCHAR2(50));
```

Example 6.2 This example creates a table, called *Customer*, with the columns: *CustId*, *CustName*, *Addr*, *Zip*, and *Region*. The *CustId* column is designated as the primary key, so that the *CustId* value in a row uniquely identifies that row in the table, as described in [“Primary keys, foreign keys and unique indexes” on page 129](#). The UNIQUE HASH ON (custId) PAGES value indicates that there are 30 pages in the hash index. This number is used to determine the number of buckets that are to be allocated for the table’s hash index. Bucket count = (PAGES * 256) / 20. Therefore the number of buckets allocated for the hash index is 384: $(30 * 256) / 20 = 384$

```
CREATE TABLE Customer
  (custId NUMBER NOT NULL PRIMARY KEY,
   custName CHAR(100) NOT NULL,
   Addr CHAR(100),
   Zip NUMBER,
   Region CHAR(10))
UNIQUE HASH ON (custId) PAGES = 30;
```

Dropping a table

To drop a TimesTen table, call the SQL statement [DROP TABLE](#).

Example 6.3 The following example drops the table *NameID*.

```
DROP TABLE NameID;
```

Estimating table size

Increasing the size of a TimesTen data store can be done on first connect. To avoid having to increase the size of a data store, it is important not to

underestimate the eventual data store size. Use the utility **ttSize** to estimate data store size.

Implementing aging in your tables

You can define an aging policy for one or more tables in your data store. An aging policy refers to the type of aging and the aging attributes, as well as the aging state (ON or OFF). You can specify one of the following types of aging policies: usage-based or time-based. Usage-based aging removes least recently used (LRU) data within a specified data store usage range. Time-based aging removes data based on the specified data lifetime and frequency of the aging process. You can define both usage-based aging and time-based aging in the same data store, but you can define only one type of aging on a specific table.

You can define an aging policy for a new table with the **CREATE TABLE** statement. You can add an aging policy to an existing table with the **ALTER TABLE** statement if the table does not already have an aging policy defined. You can change the aging policy by dropping aging and adding a new aging policy.

You cannot specify aging on the following types of tables:

- Global temporary tables
- Detail tables for materialized views

You can also implement aging in cache groups. See "[Implementing aging in a cache group](#)" in *TimesTen Cache Connect to Oracle Guide*.

This section includes the following topics:

- [Usage-based aging](#)
- [Time-based aging](#)
- [Aging and foreign keys](#)
- [Scheduling when aging starts](#)
- [Aging and replication](#)

Usage-based aging

Usage-based aging enables you to maintain the amount of memory used in a data store within a specified threshold by removing the least recently used (LRU) data.

Define LRU aging for a new table by using the AGING LRU clause of the **CREATE TABLE** statement. Aging begins automatically if the aging state is ON.

Use the **ttAgingLRUConfig** built-in procedure to specify the LRU aging attributes. The attribute values apply to all tables in the data store that have an LRU aging policy. If you do not call the **ttAgingLRUConfig** built-in procedure, then the default values for the attributes are used.

The following table summarizes the LRU aging attributes:

LRU Aging Attribute	Description
<i>LowUsageThreshhold</i>	The percent of the data store PermSize at which LRU aging is deactivated.
<i>HighUsageThreshhold</i>	The percent of the data store PermSize at which LRU aging is activated.
<i>AgingCycle</i>	The number of minutes between aging cycles.

If you set a new value for *AgingCycle* after an LRU aging policy has already been defined, aging occurs based on the current time and the new cycle time. For example, if the original aging cycle is 15 minutes and LRU aging occurred 10 minutes ago, aging is expected to occur again in 5 minutes. However, if you change the *AgingCycle* parameter to 30 minutes, then aging occurs 30 minutes from the time you call the **ttAgingLRUConfig** procedure with the new value for *AgingCycle*.

If a row has been accessed or referenced since the last aging cycle, it is not eligible for LRU aging. A row is considered to be accessed or referenced if one of the following is true:

- The row is used to build the result set of a **SELECT** statement.
- The row has been flagged to be updated or deleted.
- The row is used to build the result set of an **INSERT SELECT** statement.

You can use the **ALTER TABLE** statement to perform the following tasks:

- Enable or disable the aging state on a table that has an aging policy defined by using the **ALTER TABLE** statement with the SET AGING {ON|OFF} clause.
- Add an LRU aging policy to an existing table by using the **ALTER TABLE** statement with the ADD AGING LRU [ON|OFF] clause.
- Drop aging on a table by using the **ALTER TABLE** statement with the DROP AGING clause.

Use the **ttAgingScheduleNow** built-in procedure to schedule when aging starts. For more information, see “Scheduling when aging starts” on page 134.

To change aging from LRU to time-based on a table, first drop aging on the table by using the **ALTER TABLE** statement with the DROP AGING clause. Then add time-based aging by using the **ALTER TABLE** statement with the ADD AGING USE clause.

Note: When you drop LRU aging or add LRU aging to tables that are referenced in commands, TimesTen marks the compiled commands invalid. The commands need to be recompiled.

Time-based aging

Time-based aging removes data from a table based on the specified data lifetime and frequency of the aging process. Specify a time-based aging policy for a new table with the AGING USE clause of the [CREATE TABLE](#) statement. Add a time-based aging policy to an existing table with the ADD AGING USE clause of the [ALTER TABLE](#) statement.

The AGING USE clause has a *ColumnName* argument. *ColumnName* is the name of the column that is used for time-based aging. For brevity, we will call it the *timestamp column*. The timestamp column must be defined as follows:

- ORA_TIMESTAMP, TT_TIMESTAMP, ORA_DATE or TT_DATE data type
- NOT NULL

Your application updates the values of the timestamp column. If the value of this column is unknown for some rows and you do not want the rows to be aged, then define the column with a large default value. You can create an index on the timestamp column for better performance of the aging process.

Note: You cannot add or modify a column in an existing table and then use that column as a timestamp column because you cannot add or modify a column and define it to be NOT NULL.

You cannot drop the timestamp column from a table that has a time-based aging policy.

If the data type of the timestamp column is ORA_TIMESTAMP, TT_TIMESTAMP, or ORA_DATE, you can specify the lifetime in days, hours, or minutes in the LIFETIME clause of the [CREATE TABLE](#) statement. If the data type of the timestamp column is TT_DATE, specify the lifetime in days.

The value in the timestamp column is subtracted from SYSDATE. The result is truncated the result using the specified unit (minute, hour, day) and compared with the specified LIFETIME value. If the result is greater than the LIFETIME value, then the row is a candidate for aging.

Use the CYCLE clause to indicate how often the system should examine the rows to remove data that has exceeded the specified lifetime. If you do not specify CYCLE, aging occurs every five minutes. If you specify 0 for the cycle, then aging is continuous. Aging begins automatically if the state is ON.

Use the [ALTER TABLE](#) statement to perform the following tasks:

- Enable or disable the aging state on a table with a time-based aging policy by using the SET AGING {ON|OFF} clause.
- Change the aging cycle on a table with a time-based aging policy by using the SET AGING CYCLE clause.
- Change the lifetime by using the SET AGING LIFETIME clause.
- Add time-based aging to an existing table with no aging policy by using the ADD AGING USE clause.
- Drop aging on a table by using the DROP AGING clause.

Use the [ttAgingScheduleNow](#) built-in procedure to schedule when aging starts. For more information, see “Scheduling when aging starts” on page 134.

To change the aging policy from time-based aging to LRU aging on a table, first drop time-based aging on the table. Then add LRU aging by using the [ALTER TABLE](#) statement with the ADD AGING LRU clause.

Aging and foreign keys

Tables that are related by foreign keys must have the same aging policy.

If LRU aging is in effect and a row in a child table has been recently accessed, then neither the parent row nor the child row will be deleted.

If time-based aging is in effect and a row in a parent table is a candidate for aging out, then the parent row and all of its children will be deleted.

If a table has ON DELETE CASCADE enabled, the setting is ignored.

Scheduling when aging starts

Use the [ttAgingScheduleNow](#) built-in procedure to schedule the aging process. The aging process starts as soon as you call the procedure unless there is already an aging process in progress, in which case it will begin when that aging process has completed.

When you call [ttAgingScheduleNow](#), the aging process starts regardless of whether the state is ON or OFF.

The aging process starts only once as a result of calling [ttAgingScheduleNow](#). Calling [ttAgingScheduleNow](#) does not change the aging state. If the aging state is OFF when you call [ttAgingScheduleNow](#), then the aging process starts, but it does not continue after the process is complete. To continue aging, you must call [ttAgingScheduleNow](#) again or change the aging state to ON.

If the aging state is already set to ON, then [ttAgingScheduleNow](#) resets the aging cycle based on the time [ttAgingScheduleNow](#) was called.

You can control aging externally by disabling aging by using the [ALTER TABLE](#) statement with the SET AGING OFF clause. Then use [ttAgingScheduleNow](#) to start aging at the desired time.

Use `ttAgingScheduleNow` to start or reset aging for an individual table by specifying its name when you call the procedure. If you do not specify a table name, then `ttAgingScheduleNow` will start or reset aging on all of the tables in the data store that have aging defined.

Aging and replication

For active standby pairs, implement aging on the active master data store. Deletes that occur as a result of aging will be replicated to the standby master data store and the read-only subscribers. If a failover to the standby master data store occurs, aging is enabled on the data store after its role changes to ACTIVE.

For all other types of replication schemes, implement aging separately on each node. The aging policy must be the same on all nodes.

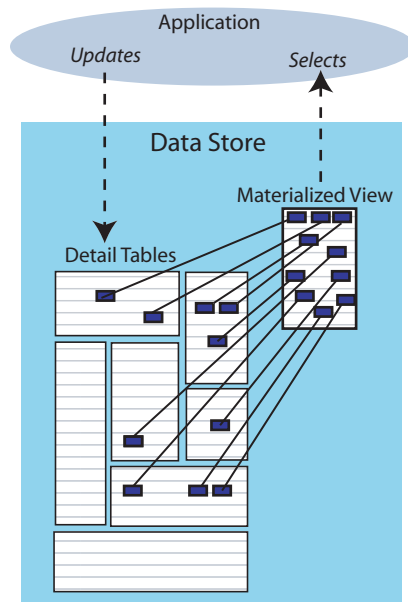
If you implement LRU aging on a multimaster replication scheme used as a hot standby, LRU aging may provide unintended results. After a failover, you may not have all of the desired data because aging occurs locally.

Understanding materialized views

A materialized view is a read-only table that maintains a summary of data selected from one or more “regular” TimesTen tables. The summary data in a materialized view is called a *result set* and the “regular” TimesTen tables queried to make up the result set are called *detail tables*.

[Figure 6.1](#) shows a materialized view created from detail tables. An application updates the detail tables and can select data from the materialized view.

Figure 6.1 Materialized View



Working with materialized views

This section includes the following topics:

- [Creating a materialized view](#)
- [Restrictions on materialized views and detail tables](#)
- [Performance implications of materialized views](#)
- [Dropping a materialized view](#)

Creating a materialized view

To create a materialized view, use the SQL statement [CREATE MATERIALIZED VIEW](#). The syntax for all SQL statements is provided in the [Oracle TimesTen In-Memory Database SQL Reference Guide](#).

Assume the following two tables have been created:

```
CREATE TABLE customer(custId int not null,  
    custName char(100) not null,  
    Addr char(100),  
    Zip int,  
    Region char(10),  
    PRIMARY KEY (custId));
```

```
CREATE TABLE bookOrder(orderId int not null,  
    custId int not null,
```

```

    ordNo int not null,
    book char(100),
    PRIMARY KEY (orderId),
    FOREIGN KEY (custId) REFERENCES Customer(custId));

```

The following creates a materialized view, named *SampleMV*, that generates a result set from selected columns in the *customer* and *bookOrder* detail tables described above.

```

CREATE MATERIALIZED VIEW SampleMV AS
    SELECT customer.custId, custName, ordNo, book
    FROM customer, bookOrder
    WHERE customer.custId=bookOrder.custId;

```

When creating a materialized view, you can establish primary keys and the size of the hash table in the same manner as described for tables in [“Primary keys, foreign keys and unique indexes” on page 129](#).

The SELECT query in the CREATE MATERIALIZED VIEW statement

The [SELECT](#) query used to define the contents of a materialized view is similar to the top-level SQL [SELECT](#) statement described in [Chapter 5, “SQL Statements”](#) in the *Oracle TimesTen In-Memory Database SQL Reference Guide*, with the following restrictions:

- All columns in the GROUP BY *GroupColumnList* must be included in the *SelectList*.
- Aggregate views must include a COUNT(*) in the *SelectList* so that TimesTen can do incremental updates of a group. For example, a group should be removed if its count becomes zero.
- SUM and COUNT are allowed, but not expressions involving them, including AVG.
- The following cannot be used in a [SELECT](#) statement that is creating a materialized view:
 - DISTINCT
 - FIRST
 - HAVING
 - ORDER BY
 - UNION
 - UNION ALL
 - MINUS
 - INTERSECT
 - JOIN
 - User functions: USER, CURRENT_USER, SESSION_USER
 - Subqueries

- NEXTVAL and CURRVAL
- Derived tables and joined tables
- Each expression in the *SelectList* must have a unique name. A name of a simple column expression would be that column's name unless a column alias is defined. ROWID is considered an expression and needs an alias.
- OUTER JOINS are allowed, but the **SELECT** list must project at least one non-nullable column from each of the inner tables specified in the OUTER JOIN. Outer join syntax for a **SELECT** in a materialized view definition is identical to that in a top-level **SELECT**. The restrictions noted in the description of **SELECT** statements apply. The (+) symbol must be used to specify OUTER JOINS of a materialized view.
- Self joins are allowed. A self join is a join of a table to itself. This table appears more than once in the FROM clause and is followed by table aliases that qualify column names in the join condition. Materialized views created with self-join in this release of TimesTen are not compatible with materialized views from releases earlier than 6.0.

Restrictions on materialized views and detail tables

After a materialized view is created, changes made to the data in the detail tables are immediately reflected in the materialized view. The only way to update a materialized view is by changing the viewed data in the detail tables. A materialized view is a read-only table that cannot be updated directly. This means a materialized view cannot be updated by an INSERT, DELETE, or UPDATE statement by replication, XLA, or the cache agent (it cannot be part of a cache group).

For example, an attempt to update a row in a materialized view generates the error:

```
805: Update view table directly has not been implemented
```

Readers familiar with other implementations of materialized views should note the following characteristics of TimesTen views:

- Detail tables can be replicated, but materialized views cannot.
- Neither a materialized view nor its detail tables can be part of a cache group.
- TimesTen does not automatically create indexes on materialized views or detail tables. No referential indexes can be defined on the materialized view.
- A materialized view cannot be dropped with a **DROP TABLE** statement. You must use the **DROP VIEW** statement.
- A materialized view cannot be altered with an **ALTER TABLE** statement. You must use the **DROP VIEW** statement and then create a new materialized view with a **CREATE MATERIALIZED VIEW** statement.

- Materialized views must be explicitly created by the application. The TimesTen query optimizer has no facility to automatically create materialized views.
- The TimesTen query optimizer does not rewrite queries on the detail tables to reference materialized views. Application queries must directly reference views, if they are to be used.
- There is no deferred maintenance option for materialized views. A materialized view is refreshed automatically when changes are made to its detail tables and there is no facility for manually refreshing a view.
- There are some restrictions to the SQL used to create materialized views. See [CREATE MATERIALIZED VIEW](#) in the *Oracle TimesTen In-Memory Database SQL Reference Guide* for details.

Performance implications of materialized views

The performance of **UPDATE** and **INSERT** operations may be impacted if the updated table is referenced in a materialized view. The performance impact depends on many factors, such as the nature of the view (how many detail tables, whether outer join or aggregation is used), which indexes are present on the detail table and on the view, and how many view rows will be affected by the change.

A view is a persistent, up-to-date copy of a query result. To keep the view up to date, TimesTen must do “view maintenance” when you change a view’s detail table. For example, if you have a view named *V* that selects from tables *T1*, *T2*, and *T3*, then any time you insert into *T1*, or update *T2*, or delete from *T3*, TimesTen does “view maintenance.”

View maintenance needs appropriate indexes just like regular database operations. If they’re not there, view maintenance will take a very long time.

All of the updates/inserts/deletes on detail tables have execution plans, as described in [Chapter 9, “The TimesTen Query Optimizer.”](#) For example, an update of a row in *T1* will have a first stage of the plan where it updates the view *V*, followed by a second stage where it updates *T1*.

For view maintenance to be fast, you should look at the plans for all the operations that update (write to) the detail tables and make sure they look good. This may be time-consuming, so here’s an easy method to start off with:

1. Look at all the WHERE clauses for the updates/deletes that frequently happen on the detail tables. Those clauses will probably use an index key. Make a note of each of those keys. For example, if the operations that an application does 95% of the time are:
 - * update T1 set A=A+1 WHERE K1=? AND K2=?
 - * delete from T2 WHERE K3=?

Then the keys you're remembering are (K1,K2) and K3.

2. Make sure that the view selects all of those key columns. In this example, the view should select K1, K2, and K3.
3. Create an index on the view on each of those keys. In this example, the view should have two indexes, one on (V.K1,V.K2) and one on V.K3. The indexes don't have to be unique, though they can be. (And the names of the view columns can be different from the names of the table columns, though they're the same in this example.)

With this method, when you update a detail table, your WHERE clause is used to do the corresponding update of the view. This allows maintenance to be done in a batch, and is very fast.

The above method may not always work, however. For example, an application may have many different ways of updating the detail tables, and so would have to select far too many things in the view, or create too many indexes on the view, taking up more space or more performance than you might like. So, if the above is too cumbersome, an alternative method is:

1. For each table in the view's FROM clause (each detail table), check which ones are frequently changed by UPDATE, INSERT and CREATE VIEW statements. For example, a view's FROM clause may have tables T1, T2, T3, T4, T5, but of those, only T2 and T3 are frequently changed.
2. For each of those tables, make sure the view selects their rowids. In this example, the view should select T2.rowid and T3.rowid.
3. Create an index on the view on each of those rowid columns. In this example, the columns might be called T2rowid and T3rowid, and indexes would be created on V.T2rowid and V.T3rowid.

With this method, view maintenance is done on a row-by-row basis, rather than on a batch basis. But the rows can be matched very efficiently between a view and its detail tables. So that speeds up the maintenance. It's generally not as fast as the first method, but it's still pretty good.

Dropping a materialized view

To drop a materialized view, execute the DROP VIEW SQL statement.

Example 6.4 The following statement drops the *sampleMV* materialized view, *sampleMV*.

```
DROP VIEW sampleMV;
```

Understanding views

A *view* is a logical table that is based on one or more tables. The view itself contains no data. It is sometimes called a *nonmaterialized view* to distinguish it from a materialized view, which does contain data that has already been

calculated from *detail tables*. Views cannot be updated directly, but changes to the data in the detail tables are immediately reflected in the view.

To choose whether to create a view or a materialized view, consider where the cost of calculation lies. For a materialized view, the cost falls on the users who update the detail tables because calculations must be made to update the data in the materialized views. For a nonmaterialized view, the cost falls on a connection that queries the view, because the calculations must be made at the time of the query.

Working with views

This section includes the following topics:

- [Creating a view](#)
- [The SELECT query in the CREATE VIEW statement](#)
- [Restrictions on views and their detail tables](#)
- [Dropping a view](#)

Creating a view

To create a view, use the [CREATE VIEW](#) SQL statement. The syntax for all SQL statements is provided in the [Oracle TimesTen In-Memory Database SQL Reference Guide](#).

```
CREATE VIEW ViewName AS SelectQuery;
```

This selects columns from the detail tables to be used in the view.

For example, create a view from the table t1:

```
CREATE VIEW v1 AS SELECT * FROM t1;
```

Now create a view from an aggregate query on the table t1:

```
CREATE VIEW v1 (max1) AS SELECT max(x1) FROM t1;
```

The SELECT query in the CREATE VIEW statement

The [SELECT](#) query used to define the contents of a materialized view is similar to the top-level SQL [SELECT](#) statement described in [Chapter 5, “SQL Statements”](#) in the [Oracle TimesTen In-Memory Database SQL Reference Guide](#), with the following restrictions:

- A [SELECT *](#) query in a view definition is expanded at view creation time. Any columns added after a view is created do not affect the view.
- The following cannot be used in a [SELECT](#) statement that is creating a view:
 - DISTINCT
 - FIRST
 - ORDER BY

- Arguments
- Temporary tables
- Each expression in the select list must have a unique name. A name of a simple column expression would be that column's name unless a column alias is defined. *RowId* is considered an expression and needs an alias.
- No `SELECT FOR UPDATE` or `SELECT FOR INSERT` statements can be used on a view.
- Certain TimesTen query restrictions are not checked when a non-materialized view is created. Views that violate those restrictions may be allowed to be created, but an error is returned when the view is referenced later in an executed statement.

Restrictions on views and their detail tables

Views have the following restrictions:

- When a view is referenced in the `FROM` clause of a `SELECT` statement, its name is replaced by its definition as a derived table at parsing time. If it is not possible to merge all clauses of a view to the same clause in the original select to form a legal query without the derived table, the content of this derived table is **materialized**. For example, if both the view and the referencing select specify aggregates, the view is **materialized** before its result can be joined with other tables of the select.
- A view cannot be dropped with a `DROP TABLE` statement. You must use the `DROP VIEW` statement.
- A view cannot be altered with an `ALTER TABLE` statement.
- Referencing a view can fail due to dropped or altered detail tables.

Dropping a view

The `DROP VIEW` statement deletes the specified view.

The following statement drops the `CustOrder` view:

```
DROP VIEW CustOrder;
```

Understanding indexes

Indexes are auxiliary data structures that greatly improve the performance of table searches. They are used automatically by the query optimizer to speed up the execution of a query. For information about the query optimizer, see [Chapter 9, “The TimesTen Query Optimizer.”](#)

TimesTen provides two types of indexes to enable fast access to tables: *hash* indexes and *T-tree* (or *range*) indexes.

- **Hash Indexes.** Hash indexes are useful for finding rows with an exact match on one or more columns. TimesTen currently supports a maximum of one hash index per table. The hash index must be over the primary key of a table. A hash index is created automatically when a table is created with a primary key specified on one or more columns or with the UNIQUE HASH option specified on one or more columns of the table.

The section “CREATE TABLE” of the [Oracle TimesTen In-Memory Database SQL Reference Guide](#) discusses in detail the automatic creation of hash indexes.

- **T-tree Indexes.** T-tree indexes are useful for finding rows with column values within a certain range. T-trees were designed specifically for in-memory applications. T-tree indexes may be created over one or more columns of a table, and up to 32 T-tree indexes may be created on one table.

Hash indexes are useful for doing equality searches, while T-trees and equijoins can be used for equality and range (greater than/equal to, less than/equal to, etc.) searches. So if you have a primary key on a field and want to see if `FIELD > 10`, the primary key index will not expedite finding the answer, but a separate index will.

Note: Hash indexes are faster than T-tree indexes for exact match lookups, but they require more space than T-tree indexes. Hash indexes cannot be used for lookups involving ranges.

You can designate an index as *unique*, which means that each row in the table has a unique value for the indexed column or columns. Unique indexes can be created over nullable columns. In conformance with the SQL standard, multiple NULL values are permitted in a unique index.

When sorting data values, TimesTen considers NULL values to be larger than all non-NULL values. See the [Oracle TimesTen In-Memory Database SQL Reference Guide](#) for more information on NULL values.

In addition to hash and T-tree indexes, lookups by RowID can be used for fast access to data. RowID lookups are faster than either type of index lookup. See the [Oracle TimesTen In-Memory Database SQL Reference Guide](#) for more information on RowIDs.

Working with indexes

This section includes the following topics:

- [Creating an index](#)
- [Altering an index](#)
- [Dropping an index](#)
- [Estimating index size](#)

Creating an index

To create an index, execute the SQL statement [CREATE INDEX](#). TimesTen converts index names to upper case characters.

Every index has an owner. The owner is the user who created the underlying table. Indexes created by TimesTen itself, such as indexes on system tables, are created with the user name `SYS` or with the user name `TTREP` if created during replication.

Example 6.5 Create an index *IxID* over column *CustID* of table *NameID*.

```
CREATE INDEX IxID ON NameID (CustID);
```

You can also create a hash index by creating a primary key or using the `UNIQUE HASH ON` clause in the [CREATE TABLE](#). However, TimesTen may create temporary hash and T-tree indexes automatically during query processing to speed up query execution.

Altering an index

You can change a hash index to a t-tree index with the `USE TREE INDEX` of the [ALTER TABLE](#) statement.

You can change a t-tree index to a hash index with the `USE HASH INDEX` of the [ALTER TABLE](#) statement.

Dropping an index

To uniquely refer to an index, an application must specify both its owner and name. If the application does not specify an owner, TimesTen looks for the index first under the user name of the caller, then under the user name `SYS`.

To drop a TimesTen index, execute the [DROP INDEX](#) SQL statement. All indexes in a table are dropped automatically when the table is dropped.

Example 6.6 The following drops the index named `IxID`.

```
DROP INDEX IxID;
```

Estimating index size

Increasing the size of a TimesTen data store can be done on first connect. To avoid having to increase the size of a data store, it is important not to underestimate the eventual data store size. Use the utility [ttSize](#) to estimate data store size.

Understanding rows

Rows are used to store TimesTen data. TimesTen supports several data types for fields in a row, including:

- One-byte, two-byte, four-byte and eight-byte integers.
- Four-byte and eight-byte floating-point numbers.
- Fixed-length and variable-length character strings, both ASCII and Unicode.
- Fixed-length and variable-length binary data.
- Fixed-length fixed-point numbers.
- Time represented as hh:mm:ss [AM|am|PM|pm].
- Date represented as yyyy-mm-dd.
- Timestamp represented as yyyy-mm-dd hh:mm:ss.

Chapter 1, “Data Types” in the *Oracle TimesTen In-Memory Database SQL Reference Guide* contains a detailed description of these data types.

Working with rows

This section includes the following topics:

- [Inserting rows](#)
- [Deleting rows](#)

Inserting rows

To insert a row, execute [INSERT](#) or [INSERT SELECT](#). You can also use the [ttBulkCp](#) utility.

Example 6.7 To insert a row in the table *NameID*, enter

```
INSERT INTO NameID VALUES(23125, 'John Smith');
```

Note: When inserting multiple rows into a table, it is more efficient to use prepared commands and parameters in your code. Create indexes after the bulk load is completed.

Deleting rows

To delete a row, execute the [DELETE](#) statement.

Example 6.8 The following deletes all the rows from the table *NameID* for names that start with the letter “S.”

```
DELETE FROM NameID WHERE CustName LIKE 'S%';
```


Transaction Management and Recovery

TimesTen supports transactions that provide atomic, consistent, isolated and durable (ACID) access to data.

TimesTen transactions support ANSI Serializable and ANSI Read_Committed levels of isolation. ANSI Serializable isolation is the most stringent transaction isolation level. ANSI Read_Committed allows greater concurrency.

Read_Committed is the default and is an appropriate isolation level for most applications. These isolation levels can be combined with each other and with a range of durability options.

TimesTen allows applications to choose the transaction features they need so they do not incur the performance overhead of features they do not need. See [Chapter 1, “Data Store Attributes”](#) in the *Oracle TimesTen In-Memory Database API Reference Guide* for details on how to set isolation levels and durability options.

The main topics in this chapter are:

- [Transaction semantics](#)
- [Transaction atomicity and durability](#)
- [Controlling durability and logging](#)
- [Concurrency control](#)
- [Checkpoints](#)

Transaction semantics

TimesTen maintains user-specified levels of isolation, atomicity and durability. As a transaction modifies data in a data store, locking, versioning and logging are used to ensure ACID properties:

- **Locking.** TimesTen acquires locks on data items that the transaction reads and/or writes, depending on the transaction isolation level. See [“Concurrency control” on page 157](#).
- **Logging.** Modifications to the data store are recorded at user-specified levels in a log. See [“Transaction atomicity and durability” on page 151](#).
- **Versioning.** TimesTen makes multiple copies of data items to allow non-serializable read and write operations on those data items to proceed in parallel.

The following table shows how TimesTen uses locks and logs:

If	Then
Transaction is terminated successfully (committed)	<ul style="list-style-type: none">• Log is posted to disk (if the DurableCommits attribute is turned on).• Newly modified values of data are made available for other transactions to read and to modify.• Locks that were acquired on behalf of the transaction are released and the corresponding data becomes available to other transactions.
Transaction is rolled back	<ul style="list-style-type: none">• Log is used to undo the effects of the transaction and to restore any modified data items to the state they were before the transaction began.• Locks that were acquired on behalf of the transaction are released and the corresponding data becomes available to other transactions.

If	Then
System fails (data not committed)	<ul style="list-style-type: none"> On first connect, TimesTen automatically performs data store recovery by reading the latest checkpoint image and applying the log to restore the data store to its latest transactionally consistent state. (See “Checkpoints” on page 160).
Application fails	<ul style="list-style-type: none"> Transaction is rolled back, if possible. If rollback is not possible, other connections to data store may be invalidated. Recovery is done on next connect (See “Recovery after fatal errors” on page 22.)

TimesTen supports temporary data stores, which have essentially no checkpoints. Recovery is never done for such data stores. They will be destroyed after a data store or application shuts down or fails.

TimesTen supports data stores with no logging. Rollback is not possible for such data stores, and recovery uses only the latest checkpoint. This mode is suitable only for special-purpose operations, such as bulk-loading a data store.

Transactions are started automatically on behalf of an application as needed. Virtually all operations on the data store, even those that do not modify or access application data, require transactional access. For example, compaction and checkpoint operations begin a transaction if one has not already been started. An application can commit a transaction by calling the ODBC **SQLTransact** (*henv*, *hdbc*, SQL_COMMIT) function or JDBC **Connection.commit** method, or abort it by calling the ODBC **SQLTransact** (*henv*, *hdbc*, SQL_ROLLBACK) function or **Connection.rollback** method. Any subsequent data store operation will automatically cause a new transaction to be started.

In compliance with ODBC standards, the default AUTOCOMMIT setting is ON. Commits are costly for performance and can be intrusive if they are implicit. TimesTen recommends that you turn AUTOCOMMIT off so that commits are intentional. Use the ODBC **SQLSetConnectOption** function or JDBC **Connection.setAutoCommit**(false) method in your TimesTen application to set SQL_AUTOCOMMIT_OFF.

When using ODBC or JDBC batch operations to **INSERT**, **UPDATE** or **CREATE VIEW** several rows in one call, when **AUTOCOMMIT** is on, a commit occurs after the entire batch operation has completed. If there is an error during the batch operation, those rows that have been successfully modified will be committed. If an error occurs due to a problem on a particular row, the preceding rows are committed. The **pirow** parameter to the ODBC **SQLParamOptions** function contains the number of the row in the batch that had the problem.

Even with durable commits and autocommit enabled, you could lose work if there is a failure or the application exits without closing cursors. An open cursor under **AUTOCOMMIT** means that you are in effect running with **AUTOCOMMIT** off but without the ability to rollback. Write locks (from DDL or DML) are held until all cursors are closed.

Note: Autocommit is the default mode for ODBC applications. Applications must explicitly turn autocommit off to avoid it.

Transaction atomicity and durability

The TimesTen Data Manager provides durability with a combination of checkpointing and logging. (See “Checkpoints” on page 160 and “Log files” on page 156.)

Because transaction support adds overhead to execution time, TimesTen allows applications to choose from the following options:

- *Guaranteed atomicity and durability.* (**Logging=1, DurableCommits=1**)
- *Guaranteed atomicity, delayed durability.* (**Logging=1, DurableCommits=0**)
- *No guaranteed atomicity, no guaranteed durability.* (**Logging=0, DurableCommits=0**)

Overview

The following table summarizes the guarantees and limitations of the atomicity and durability options:

Logging Attribute Setting	Non-blocking checkpoints possible?	Row-level locking possible?	Transaction rollback possible?	Recovery procedure	Transactions vulnerable to loss
Logging = 1 Durable Commits = 1	Yes	Yes	Yes	Read most recent checkpoint image. Apply log.	Uncommitted transactions at the time of failure
Logging = 1 Durable Commits = 0	Yes	Yes	Yes	Read most recent checkpoint image. Apply log.	Transactions that committed after the last checkpoint or durable commit

Logging Attribute Setting	Non-blocking checkpoints possible?	Row-level locking possible ?	Transaction rollback possible?	Recovery procedure	Transactions vulnerable to loss
Logging = 2 Durable Commits = 0	No	Yes	Yes	Read most recent checkpoint image. There is no log to apply.	Transactions that committed after the last checkpoint
Logging = 0 Durable Commits = 0	No	No	No	Read most recent checkpoint image. There is no log to apply.	Transactions that committed after the last checkpoint

The sections that follow description these options in greater detail.

Guaranteed atomicity and durability

(**Logging=1, DurableCommits=1**)

By default, all TimesTen transactions are durable. The effects of the transaction are persistent and will not be lost in the event of system failure.

Durability is implemented with a combination of checkpointing and logging. (See “[Checkpoints](#)” on page 160 and “[Log files](#)” on page 156.) A checkpoint operation writes the current data store image to a checkpoint file on disk, which has the effect of making all transactions that committed before the checkpoint durable. For transactions that committed after the last checkpoint, TimesTen uses conventional logging techniques to make them durable. As each transaction progresses, it records its data store modifications in an in-memory log. At commit time, the relevant portion of the log is flushed to disk. This log flush operation makes that transaction, and all previously-committed transactions, durable.

In the case of a system failure, recovery uses the last checkpoint image together with the log to reconstruct the latest transaction-consistent state of the data store.

In addition to being durable, by default all TimesTen transactions are also atomic. Either all or none of the effects of the transaction is applied to the data store.

Atomicity is implemented by using the log to undo the effects of a transaction if it is rolled back. Rollback can be caused explicitly by the application, using the

ODBC **SQLTransact** function or JDBC **Connection.rollback** method, or during data store recovery because the transaction was not committed at the time of failure.

In order to have guaranteed atomicity and durability, applications must set the **Logging** and **DurableCommits** attributes to **1**.

Guaranteed atomicity, delayed durability

(**Logging=1, DurableCommits=0**)

It is possible to connect to a data store with logging enabled but with guaranteed durability disabled. In this case, the atomicity of a transaction is guaranteed, but not its durability. This mode is known as delayed durability mode.

In delayed durability mode, as in guaranteed durability mode, each transaction enters records into the in-memory log as it makes modifications to the data store. However, when a transaction commits in delayed durability mode, it does not wait for the log to be posted to disk before returning control to the application. Since the content of the in-memory log would be lost in a system failure, transactions committed in this mode are *not* durable.

Non-durably committed transactions typically have better response time than durably-committed transactions, because no I/O is required to commit the transaction. A non-durably committed transaction can be made durable either by checkpointing (See “[Checkpoints](#)” on page 160.) or by committing a subsequent transaction durably. As with guaranteed durability mode, a checkpoint makes all transactions that committed before the checkpoint durable. Committing a transaction durably commits that transaction and all previously committed transactions.

Applications that wish to take advantage of the performance benefits of delayed durability mode but which can only tolerate the loss of a small number of transactions can perform periodic durable commits in a background process — only those transactions that committed non-durably after the last durable commit are vulnerable to loss in the event of a system failure.

Applications request delayed durability mode by setting the **Logging** attribute to **1** (the default) and the **DurableCommits** attribute to **0**. When in this mode, applications can call the **ttDurableCommit** built-in procedure to force the current transaction to commit durably when it commits.

No guaranteed atomicity, no guaranteed durability

(**Logging=0, DurableCommits=0**)

For the best response time, applications can turn off logging altogether. However, by doing so, they risk both a loss of durability and atomicity. Therefore, the no-logging mode should be used with care. (See “[Logging](#)” in the *Oracle TimesTen In-Memory Database API Reference Guide*.)

The only way to make a transaction durable when logging is disabled is to checkpoint the data store. Therefore, a system failure will result in the loss of all transactions that committed after the last checkpoint. In addition, some operations that would fail atomically with logging enabled will fail non-atomically with logging disabled. For example, a SQL **UPDATE** statement that encounters an error after already having updated one or more rows will not restore the original contents of the updated rows. When this situation arises, a warning is returned indicating that the operation has failed non-atomically.

With logging disabled, transaction rollback is not possible as there is no log of the data store modifications made by the transaction. Therefore all transactions must commit. Attempts to roll a transaction back in no-logging mode will result in an error. Row-level locking is also disabled when logging is disabled.

Applications must use data store-level locking.

Because of the many restrictions that exist in non-logging mode, its use should be limited to applications that can be restarted in the event of failure (loading data into a new data store, for example). Otherwise uncommitted data may persist, resulting in inconsistencies.

To use the no-logging mode, applications set **Logging=0**, which requires the following additional attribute settings: **LockLevel=1**, **DurableCommits=0**, and **LogPurge=0**.

Controlling durability and logging

Applications can control whether a transaction is durable and whether log files are created:

- **Durability.** By default, transactions are durable (`DurableCommits=1`). To turn off guaranteed durability, turn off the **DurableCommits** connection attribute.

A running application can override the delayed transaction durability it set for its connection by invoking the TimesTen built-in procedure **ttDurableCommit** to ensure the durability of a specific transaction.

For a shared data store, durable connections can coexist with connections that are not guaranteed durable.

- **Logging.** By default, transaction logs go to disk (`Logging=1`). To disable logging completely, set the **Logging** connection attribute to 0. An application can switch between logging to disk and no logging on the same data store after existing connections have been terminated.

All concurrent connections to a data store must have the same **Logging** attribute setting. A request for a connection whose logging attributes are incompatible with existing connections is rejected unless **MatchLogOpts** is set.

When using row-level locking or caching Oracle tables, you must use logging to disk.

Using durable commits

The performance cost of durable commits can be mitigated if many threads are running at the same time, due to an effect called “group commit.” Under group commit, a single disk write commits a group of concurrent transactions durably. Group commit does not improve the response time of any given commit operation, as each durable commit must wait for a disk write to complete, but it can significantly improve the throughput of a series of concurrent transactions.

When durable commits are used frequently, TimesTen can support more connections than there are CPUs, as long as transactions are short. This is true because each connection spends more time waiting to commit than it spends using the CPU. This is in contrast to applications that do infrequent durable commits, in which case each connection tends to be very CPU-intensive for the TimesTen portion of its workload. In the latter case, using more connections than there are processors will tend to give worse performance, due to CPU contention.

Applications that require lower response time and can tolerate some transaction loss may elect to perform periodic durable commits. By committing only every n th transaction durably, or performing a durable commit every n seconds (typically in a background process), an application can achieve quick response time while maintaining a small window of vulnerability to transaction loss.

Because a durable commit commits not only itself but all previously-committed transactions durably — even those performed by other threads or other processes, an application that commits every n transactions durably only risks the loss of the last n transactions.

Similarly, an application that performs a durable commit every n seconds, risks only the transactions that committed during the last n seconds.

To enable periodic durable commits, an application connects with **Logging=1** and **DurableCommits=0**, so transactions commit non-durably by default. When a durable commit is needed, the application calls the **ttDurableCommit** built-in procedure before committing. As with all SQL statements, it is best to pre-prepare the call to **ttDurableCommit** if it will be used frequently. The **ttDurableCommit** built-in procedure does not actually commit the transaction; it merely causes the commit to be durable when it occurs.

Another option for avoiding data loss is to use TimesTen replication instead of durable commits to maintain a copy of the data in two memories. Although two memories are not as durable as disk, replication can provide higher data availability by allowing for failover without data store recovery. This type of trade-off is common in high-performance systems. For more details on TimesTen replication, see the *TimesTen to TimesTen Replication Guide*.

Log files

If logging to disk is enabled, log files are created in the same directory as the data store files unless the **LogDir** attribute specifies a different location. Multiple log files can exist for a single data store. The log file names have the form `ds_name.logn`, where `ds_name` is the data store path name given in the data store's DSN and n is the log file number, starting at zero.

To retain archived log files, set the **LogPurge** attribute to 0. When the **LogPurge** attribute is not set for the data store connection, TimesTen renames log files no longer needed to perform recovery to `ds_name.logn.arch`. In this case, the application is responsible for removing these unneeded log files. See “**LogPurge**” in *Oracle TimesTen In-Memory Database API Reference Guide* for details on purging log files.

Concurrency control

Transaction isolation allows each active transaction to operate as if there were no other transactions active in the system. TimesTen supports two transaction isolation levels: ANSI Serializable and ANSI Read_Committed isolation.

Transaction isolation levels

- In **ANSI Serializable** isolation, each transaction acquires locks on all data items that it reads or writes. It holds these locks until it commits or rolls back. As a result, a row that has been read by one transaction cannot be updated or deleted by another transaction until the original transaction terminates. Similarly, a row that has been inserted, updated or deleted by one transaction cannot be accessed in any way by another transaction until the original transaction terminates.

Repeatable reads are assured: a transaction that executes the same query multiple times is guaranteed to see the same result set each time. Other transactions cannot **UPDATE** or **CREATE VIEW** any of the returned rows, nor can they **INSERT** a new row that satisfies the query predicate.

In this isolation level, readers can block writers and writers can block readers and other writers.

- In **ANSI Read Committed** isolation, each transaction acquires locks only on the items that it writes. Items read by **SELECT** statements, the **SELECT** portion of **INSERT SELECT** statements and **MERGE** statements, are not locked. This is the default isolation level for TimesTen.

In this isolation level, readers do not block writers, nor do writers block readers, even when they read and write the same data items. To allow readers and writers to access the same items without blocking, writers create private copies of the items that they update. These private copies become public when the transaction commits, or are discarded if the transaction rolls back.

Therefore, when a transaction reads an item that has been updated by another in-progress transaction, it sees the state of that item before it was updated. It cannot see an uncommitted state.

Non-repeatable reads are possible in this isolation level. If a read committed transaction executes the same query multiple times, the commit of an updater transaction may cause it to see different results.

Using read committed isolation level can lead to duplicates in a result set. A **SELECT** statement selects more or less rows than the total number of rows in the table if some rows were added or removed and committed in the range that the **SELECT** scan is happening. This may happen when an **UPDATE**, **INSERT** or **DELETE** adds or deletes a value from an index and the **SELECT** scan is using this index. This can also happen when an **INSERT** or **DELETE**

adds or deletes rows from the table and the SELECT operation is using an all-table scan.

This happens because index values are ordered and an UPDATE of an index value may delete the old value and insert the new value into a different place. In other words it moves a row from one position in the index to another position. If an index scan sees the same row in both positions, it returns the same row twice. This does not happen with a serial scan because table pages are unordered and rows do not need to be moved around for an UPDATE. Hence once a scan passes a row, it will not see that same row again.

The only general way to avoid this problem is for the SELECT to use serializable isolation. This prevents a concurrent INSERT, DELETE or UPDATE operation. There is no reliable way to avoid this problem with INSERT or DELETE by forcing the use of an index because these operations affect all indexes. With UPDATE, this problem can be avoided by forcing the SELECT statement to use an index that is not being updated.

All data access in TimesTen uses locking or copying to provide isolation. Applications set the transaction isolation level either using the SQLSetConnectOption ODBC function with the SQL_TXN_ISOLATION flag, or by setting the **Isolation** connection attribute.

To ensure that materialized views are always in a consistent state, all view maintenance operations are effectively performed under Serializable isolation, even when the transaction is otherwise in Read_Committed isolation. This means that the transaction will obtain read locks for any data items read during view maintenance. However, the transaction will release these read locks at the end of the **INSERT**, **UPDATE** or **CREATE VIEW** statement that triggered the view maintenance instead of holding them until it terminates.

Locking granularities

TimesTen supports row-level locking and data store-level locking:

- With row-level locking, transactions usually obtain locks on the individual rows that they access, although a transaction may obtain a lock on an entire table if TimesTen determines that doing so would result in better performance. Row-level locking is the default and is the best choice for most applications, as it provides the finest granularity of concurrency control. Row-level locking requires disk logging. Applications can control the circumstances under which TimesTen elects to use a table lock by setting optimizer flags. (See “[ttOptSetFlag](#)” in the *Oracle TimesTen In-Memory Database API Reference Guide*.) To use row-level locking, applications must set the **LockLevel** connection attribute to **0** (the default value) or call the **ttLockLevel** built-in procedure with the *lockLevel* parameter set to **Row**.
- With data store-level locking, every transaction obtains an exclusive lock on the entire data store, thus ensuring that there is no more than one active transaction in the data store at any given time. Data store-level locking often provides better performance than row-level locking, due to reduced locking overhead. However, its applicability is limited to those applications that never need to execute multiple concurrent transactions. With data store-level locking, every transaction effectively runs in ANSI Serializable isolation, since concurrent transactions are disallowed. Data store level is required when logging is disabled. To use data store-level locking, applications set the **LockLevel** connection attribute to **1** or call the **ttLockLevel** built-in procedure with the *lockLevel* parameter set to **DS**.

Coexistence of different locking levels

Different connections can coexist with different levels of locking, but the presence of even one connection using data store-level locking leads to reduced concurrency. For performance information, see “[Choose the best method of locking](#)” in the *Oracle TimesTen In-Memory Database Java Developer’s and Reference Guide*. For information on tuning TimesTen C applications, see “[Choose the best method of locking](#)” in the *Oracle TimesTen In-Memory Database C Developer’s and Reference Guide*.

Checkpoints

A checkpoint is an operation that saves the state of a data store to disk files, known as checkpoint files. By default, TimesTen performs “background” checkpoints at regular intervals. Alternatively, applications can programmatically initiate checkpoint operations. See [“Setting and managing checkpoints” on page 162](#), for more details.

Each TimesTen data store has two checkpoint files, named *dsname.ds0* and *dsname.ds1*, where *dsname* is the data store path name specified in the data store's DSN. A checkpoint operation identifies the checkpoint file to which the last checkpoint was written and writes its checkpoint to the other file. Therefore, the two files always contain the two most recent data store images.

Data store recovery uses these files to recover the most recent transaction-consistent state of the data store after a data store shutdown or system failure. It identifies the file that contains the more recent of the two checkpoint images and applies the log to that file's data store image, as appropriate, to recover the up-to-date data store state. If any errors occur during this process, or if the more recent checkpoint image is incomplete (for example, if a system failure occurred while that checkpoint was begin written), then recovery restarts, using the other checkpoint file.

TimesTen also creates a *dsName.res0* and *dsName.res1* file for each data store. These files are used internally by TimesTen for the creation of logs.

A checkpoint operation has two primary purposes. First, it decreases the amount of time required for data store recovery, because it provides a more up-to-date data store image for recovery to begin with. Second, it makes a portion of the log unneeded for any future data store recovery operation, typically allowing one or more log files to be deleted. Both of these functions are very important to TimesTen applications. The reduction in recovery time is important, as the amount of log needed to recover a data store has a direct impact on the amount of downtime seen by an application after a system failure. The removal of unneeded log files is important because it frees disk space that can then be used for new log files. If these files were never removed, they would eventually consume all available space in the log directory's file system, causing data store operations to fail due to log space exhaustion.

For these reasons, either TimesTen applications should checkpoint their data stores periodically or you should set the data store first connection attributes **CkptFrequency** or **CkptLogVolume**, which determine how often TimesTen performs a “background” checkpoint.

Checkpointing may generate a large amount of I/O activity and have a long execution time depending on the size of the data store and the number of data store changes since the most recent checkpoint.

Types of checkpoints

TimesTen supports two kinds of data store checkpoints:

Transaction-consistent checkpoints

Transaction-consistent checkpoints, also known as blocking checkpoints, obtain an exclusive lock on the data store for a portion of the checkpoint, blocking all access to the data store during that time. The resulting checkpoint image contains the effects of all transactions that committed before the checkpointer obtained its lock. Because no transactions can be active while the data store lock is held, no modifications made by in-progress transactions are included in the checkpoint image.

Transaction-consistent checkpoints can be used with any of the logging modes (disk logging and no logging). When disk logging is in effect, the log is used during recovery to reapply the effects of transactions that committed durably after the checkpoint completed. To request a transaction-consistent checkpoint, an application uses the **ttCkptBlocking** built-in procedure. The actual checkpoint is delayed until the requesting transaction commits or rolls back. If a transaction-consistent checkpoint is requested for a data store for which both checkpoint files are already up to date then the checkpoint request is ignored.

Fuzzy or non-blocking checkpoints

Fuzzy checkpoints, or non-blocking checkpoints, allow transactions to execute against the data store while the checkpoint is in progress. Fuzzy checkpoints do not obtain locks of any kind, and therefore have a minimal impact on other data store activity. Because these other transactions may modify the data store while it is being written to the checkpoint file, the resulting checkpoint image may contain some effects of transactions that were active while the checkpoint was in progress. Furthermore, different portions of the checkpoint image may reflect different points in time. For example, one portion may have been written before a given transaction committed, while another portion was written afterward. The term “fuzzy checkpoint” derives its name from this fuzzy state of the data store image. TimesTen background checkpoints are always non-blocking.

To recover from a fuzzy checkpoint, TimesTen uses the log both to bring the various portions of the checkpoint into a consistent state with one another and to reapply the effects of transactions that committed durably after the checkpoint completed. For this reason, fuzzy checkpoints can only be used when disk logging is in effect. To request a fuzzy checkpoint, an application uses the **ttCkpt** built-in procedure. If disk logging is in effect, this procedure issues a fuzzy checkpoint. If logging is disabled, then a blocking (transaction-consistent) checkpoint is requested. As with all blocking checkpoints, the actual checkpoint is delayed until the requesting transaction commits or rolls back.

Setting and managing checkpoints

By default, TimesTen performs automatic non-blocking checkpoints in the background if **Logging**=1. In this case, background checkpoints are non-blocking. (See “[Fuzzy or non-blocking checkpoints](#)” on page 161.)

Several data store attributes and built-in procedures are available to set, manage and monitor checkpoints. These include:

- **CkptFrequency** attribute
- **CkptLogVolume** attribute
- **CkptRate** attribute
- **ttCkpt** built-in procedure
- **ttCkptBlocking** built-in procedure
- **ttCkptConfig** built-in procedure
- **ttCkptHistory** built-in procedure

TimesTen also automatically performs a transaction-consistent checkpoint when the last application disconnects from the data store, unless the RAM policy is **always**. For temporary data stores with **Logging**=1, checkpoints are still taken to purge the log files. For non-disk logging Temporary data stores, background checkpointing is off. See “[Transaction-consistent checkpoints](#)”.

In addition, applications can programatically perform a checkpoint using the **ttCkpt** or **ttCkptBlocking** built in procedure. For details on how to call the **ttCkpt** and other TimesTen built-in procedures from a C or Java program, see “[Calling TimesTen built-in procedures within C applications](#)” in the *Oracle TimesTen In-Memory Database C Developer's and Reference Guide* or “[Calling TimesTen built-in procedures](#)” in the *Oracle TimesTen In-Memory Database Java Developer's and Reference Guide*.

By default, TimesTen performs background checkpoints at regular intervals. If an application attempts to perform a checkpoint while a background checkpoint is in progress, TimesTen returns an error to the application. To turn off background checkpointing, set **CkptFrequency**=0 and **CkptLogVolume**=0. You can also use the built-in procedure **ttCkptConfig** to configure background checkpointing or turn it off. The values set by **ttCkptConfig** take precedence over those set with the data store attributes.

Using these attributes and the built-in procedure, you can configure TimesTen to checkpoint either when the log files contain a certain amount of data or at a specific frequency. For information on default values and usage, see the *Oracle TimesTen In-Memory Database API Reference Guide*.

If the application attempts to back up a data store while a background checkpoint is in process, TimesTen waits until the checkpoint finishes and before beginning the backup. If a background checkpoint starts while a backup is in progress, the background checkpoint will not take place until the backup has completed. If a

background checkpoint starts while an application-initiated checkpoint is in progress, then an error results.

You can use, the **ttCkptHistory** built-in procedure to display the history of last eight checkpoints, the settings for checkpoint frequency and log volume and the status of in-progress checkpoint disk writes.

Setting the checkpoint rate for background checkpoints

By default, there is no limit to the rate at which checkpoints are written to disk. You can use the **CkptRate** attribute or the **ttCkptConfig** built-in procedure to set the maximum rate at which background checkpoints are written to disk, if you would like to have control over the rate. The rate is expressed in MB per second. Checkpoints taken during recovery and final checkpoints do not honor this rate; their rate is unlimited.

See the *Oracle TimesTen In-Memory Database API Reference Guide* for details on using these features.

Setting a rate too low can cause checkpoints to take an excessive amount of time and cause the following problems;

- Delay the purging of unneeded log files
- Delay the start of backup operations
- Increase recovery time.

When choosing a rate, you should take into consideration the amount of data written by a typical checkpoint and the amount of time checkpoints usually take. Both of these pieces of information are available through the **ttCkptHistory** built-in procedure.

In addition, you can monitor the progress of a running checkpoint by looking at the *Percent_Complete* column of the **ttCkptHistory** result set. If a running checkpoint appears to be progressing too slowly, the rate can be increased by calling the **ttCkptConfig** built-in procedure. If a call to **ttCkptConfig** changes the rate, the new rate takes effect immediately, affecting even the running checkpoint.

A simple method of calculating the checkpoint rate is:

1. Call the **ttCkptHistory** built-in procedure.
2. For any given checkpoint, subtract the *starttime* from the *endtime*.
3. Divide the number of bytes written by this elapsed time in seconds to get the number of bytes per second.
4. Divide this number by 1024*1024 to get the number of megabytes per second.

When setting the checkpoint rate, some other things to consider are:

- The specified checkpoint rate is only approximate. The actual rate of the checkpoint may be below the specified rate, depending on the hardware, system load and other factors.
- Calculating the actual checkpoint rate using the above method may produce a result that is below the requested rate. This is because the *starttime/endtime* interval includes other checkpoint activities in addition to the writing of dirty blocks to the checkpoint file.
- The *Percent_Complete* field of the **ttCkptHistory** call may show 100 percent before the checkpoint is actually complete. The *Percent_Complete* field shows only the progress of the writing of dirty blocks and does not include additional bookkeeping at the end of the checkpoint.
- When adjusting the checkpoint rate, you may also need to adjust the checkpoint frequency, as a slower rate makes checkpoints take longer, which effectively increases the minimum time between checkpoint beginnings.

Data Store Performance Tuning

An application using the TimesTen Data Manager should obtain an order of magnitude performance improvement in its data access over an application using a traditional DBMS. However, poor application design and tuning can erode the TimesTen advantage. This chapter discusses factors that can affect the performance of a TimesTen application. These factors range from subtle, such as data conversions, to more overt, such as preparing a command at each execution. This chapter explains the full range of these factors, with a section on each factor indicating:

- How the problem can be detected
- Whether the potential performance impact is large, medium, small or variable.
- Where the performance gain may be
- What the tradeoffs are

As discussed throughout this chapter, many performance problems can be identified by examining the [SYS.MONITOR](#) table.

Topics are:

- [System and data store tuning](#)
- [Client/Server tuning](#)
- [SQL tuning](#)
- [Improving performance of materialized views](#)
- [Scaling to Multiple CPUs](#)
- [XLA acknowledgement modes](#)

For information on tuning TimesTen Java applications, see [Chapter 4](#), “[Application Tuning](#)” in the *Oracle TimesTen In-Memory Database Java Developer’s and Reference Guide*. For information on tuning TimesTen C applications, see [Chapter 5](#), “[Application Tuning](#)” in the *Oracle TimesTen In-Memory Database C Developer’s and Reference Guide*.

System and data store tuning

Provide enough memory

Performance impact: Large

Configure your system so that the entire data store fits in main memory. The use of virtual memory substantially decreases performance. You will know that the data store (or working set) does not fit if a performance monitoring tool shows excessive paging or virtual memory activity.

You may have to add physical memory or configure the system software to allow a large amount of shared memory to be allocated to your process(es). TimesTen includes the [ttSize](#) utility to help you estimate the size of your data store.

Size your data store correctly

Performance impact: Variable

When you create a data store, you are required to specify a data store size. Specifically, you specify a size for the permanent partition of the data store and a size for the temporary partition of the data store.

Once you have created a data store, you can increase the size of the permanent partition of the data store and you can either increase or decrease the temporary partition of the data store. See [“Specifying the size of a data store” on page 29](#).

To make size estimates, use the [ttSize](#) utility or run the application until you can make a reasonable estimate.

Another method for estimating size requirements is to use the [SYS.MONITOR](#) table. The columns [PERM_ALLOCATED_SIZE](#), [TEMP_ALLOCATED_SIZE](#), [PERM_IN_USE_SIZE](#) and [TEMP_IN_USE_SIZE](#) show (in KB units) the currently allocated size of the data store, and the in-use size of the data store. The system updates this information each time a connection is made or released and each time a transaction is committed or rolled back.

You can monitor block-level fragmentation in the data store by using [ttBlockInfo](#).

Increase LogBufferSize if needed

Performance impact: Large

Increasing the value of [LogBufferSize](#) can have a substantial positive performance impact. If [LOG_BUFFER_WAITS](#) is increasing, then increase the value of [LogBufferSize](#).

The trade-off is that more transactions are buffered in memory and may be lost if the process crashes. If [DurableCommits](#) are enabled, increasing the default [LogBufferSize](#) value does not improve performance.

Use temporary data stores if appropriate

Performance impact: Variable

A TimesTen data store may be permanent or **Temporary**. A temporary data store disappears when the last connection goes away or when there is a system or application failure.

If you do not need to save the data store to disk, you can save checkpoint overhead by creating a temporary data store.

Temporary data stores are never fully checkpointed to disk, although the log is periodically written to disk when logging to disk is enabled. However, the amount of data written to the log for temporary data stores is less than that written for permanent data stores, allowing better performance for temporary data stores. Checkpoint operations can have significant overhead for permanent data stores, depending on data store size and activity, but have very little impact for temporary data stores. Checkpoints are still necessary to remove log files.

Avoid connection overhead

Performance impact: Large

By default, TimesTen loads an idle data store (a data store with no connections) into memory when a first connection is made to it. When the final application disconnects from a data store, a delay occurs when the data store is written to disk. If applications are continually connecting and disconnecting from a data store, the data store may be loaded and unloaded to/from memory continuously, resulting in excessive disk I/O and poor performance. Similarly, if you are using a very large data store you may want to pre-load the data store into memory before any applications attempt to use it.

To avoid the latency of loading a data store into memory, you can change the RAM policy of the data store to allow data stores to always remain in memory. The trade-off is that since the data store is never unloaded from memory, a final disconnect checkpoint never occurs. So, applications should checkpoint the data store explicitly in order to reduce the disk space taken up by log files.

Alternatively, you can specify that the data store remain in memory for a specified interval of time and accept new connections. If no new connections occur in this interval, TimesTen unloads the data store from memory and checkpoints it. You can also specify a setting to allow a system administrator to load and unload the data store from memory manually.

To change the RAM policy of a data store, use the **ttAdmin** utility.

Load the data store into RAM when duplicating

Performance impact: Large

When you duplicate a data store, use the `-ramLoad` option of the **ttAdmin** utility. This places the data store in memory, available for connections, instead of unloading it with a blocking checkpoint. See [“Avoid connection overhead” on page 167](#).

Avoid OS paging at load time

Performance impact: Medium

All of the TimesTen platform operating systems implement a dynamic file system buffer pool in main memory. If this buffer pool is allowed to be large, TimesTen and the operating system both retain a copy of the file in memory, causing some of the TimesTen shared segment to be paged out.

This behavior may not occur for data stores that are less than half of the installed memory size. On some systems, it is possible to limit the amount of main memory used by the file system. On other systems, this effect is less pronounced. On HP-UX, Solaris and Linux systems, consider using the **MemoryLock** attribute to specify whether the data store should be locked in memory. If used, the data store cannot be paged out.

On HP-UX, consider the settings for the kernel parameters `dbc_min_pct` and `dbc_max_pct`. These parameters control the minimum and maximum percent of real memory devoted to the file system. The default maximum is 50 percent. TimesTen recommends reducing the maximum to 10 percent.

Consider special options for maintenance

Performance impact: Medium

During special operations such as initial loading, you can choose different options than you would use during normal operations. In particular, consider turning **Logging** off and using data store-level locking for bulk loading; an example would be using **ttBulkCp** or **ttMigrate**. These choices can improve loading performance by a factor of two.

Check your driver

Performance impact: Large

There are two versions of the TimesTen Data Manager driver for each platform, a debugging version and a production version. Unless you are actually debugging, use the production version. The debugging library can be significantly slower. See [“Specify the DSN” on page 22](#) and [“Specify the ODBC driver” on page 17](#) for a description of the TimesTen Data Manager drivers for the different platforms.



On Windows, make sure that applications that use the ODBC driver manager use a DSN that accesses the correct TimesTen driver. Make sure that direct-linked

applications are linked with the correct TimesTen driver. An application can call the ODBC **SQLGetInfo** function with the **SQL_DRIVER_NAME** argument to determine which driver it is using.

Enable tracing only as needed

Performance impact: Large

Both ODBC and JDBC provide a trace facility to help debug applications. For performance runs, make sure that tracing is disabled except when debugging applications.

To turn the JDBC tracing on, use:

```
DriverManager.setLogStream method:  
DriverManager.setLogStream(new PrintStream(System.out, true));
```

By default tracing is off. You must call this method before you load a JDBC driver. Once you turn the tracing on, you cannot turn it off for the entire execution of the application.

Investigate alternative JVMs

Performance impact: Variable

JRockit, IBM and HP provide JVMs that may perform better than the Sun JVM.

Adjust log buffer size and CPU for a large number of subscribers

Performance impact: Large

If you are planning a replication scheme that includes a large number of subscribers, then ensure the following:

- The setting for **LogBufferSize** should result in the value of **LOG_FS_READS** in the **SYS.MONITOR** table being 0 or close to 0. This ensures that the replication agent does not have to read any log records from disk. If the value of **LOG_FS_READS** is increasing, then increase the log buffer size.
- CPU resources are adequate. The replication agent on the master data store will spawn a thread for every subscriber data store. Each thread reads and processes the log independently and needs adequate CPU resources to make progress.

Migrating data with character set conversions

Performance impact: Variable

If character set conversion is requested when migrating data stores, performance may be slower than if character set conversion is not requested.

Client/Server tuning

Work locally when possible

Performance impact: Large

Using TimesTen Client to access data stores on a remote server machine adds network overhead to your connections. Whenever possible, write your applications to access the TimesTen Data Manager locally, and link the application directly with the TimesTen Data Manager.

Use shared memory segment as IPC when client and server are on the same machine

Performance impact: Variable

The TimesTen Client normally communicates with TimesTen Server using TCP/IP sockets. If both the TimesTen Client and TimesTen Server are on the same machine, client applications show improved performance by using a shared memory segment or a UNIX domain socket for inter-process communication (IPC).

To use a shared memory segment as IPC, you must set the server options in the `tendaemon.options` file. For a description of the server options, see [“Modifying the TimesTen web server options”](#) in [Chapter 3](#) of the *Oracle TimesTen In-Memory Database Operations Guide*.

In addition, applications that use shared memory for IPC must use a logical server name (for the Client DSN) with `ttShmHost` as the Network Address. For more information, see [“Creating and configuring Client DSNs on UNIX”](#) on [page 51](#).

This feature may require a significant amount of shared memory. The TimesTen Server pre-allocates the shared memory segment irrespective of the number of existing connections or the number of statements within all connections.



If your application is running on a UNIX machine and you are concerned about memory usage, the applications using TimesTen Client ODBC driver may improve the performance by using UNIX domain sockets for communication. The performance improvement when using UNIX domain sockets is not as large as when using `ShmIPC`.

Applications that take advantage of UNIX domain sockets for local connections must use a logical server name (for the Client DSN) with `ttLocalHost` as the Network Address. For more information, see [“Creating and configuring Client DSNs on UNIX”](#) on [page 51](#). In addition, make sure that your system kernel parameters are configured to allow the number of connections you require. See [“Installation prerequisites”](#) in the *Oracle TimesTen In-Memory Database Installation Guide*.

Enable TT_PREFETCH_CLOSE for serializable transactions

Performance impact: Variable

Enable TT_PREFETCH_CLOSE for serializable transactions in client/server applications. In serializable **Isolation** mode, all transactions should be committed when executed, including read-only transactions. When TT_PREFETCH_CLOSE is enabled, the server closes the cursor and commits the transaction after the server has fetched the entire result set for a read-only query. The client should still call **SQLFreeStmt(SQL_CLOSE)** and **SQLTransact(SQL_COMMIT)**, but the calls are executed in the client and do not require a network round trip between the client and server. TT_PREFETCH_CLOSE enhances performance by decreasing the network traffic between client and server.

Do not use multiple statement handles when TT_PREFETCH_CLOSE is enabled. The server may fetch all of the result set, commit the transaction, and close the statement handle before the client is finished, resulting in the closing of all statement handles.

The following examples show how to use the TT_PREFETCH_CLOSE option with ODBC and JDBC.

Example 8.1 You can set TT_PREFETCH_CLOSE with the **SQLSetConnectOption** or **SQLSetStmtOption** ODBC functions.

```
SQLSetConnectOption (hdbc, TT_PREFETCH_CLOSE,
                    TT_PREFETCH_CLOSE_ON);
SQLExecDirect (hstmt, "SELECT * FROM T", SQL_NTS);
while (SQLFetch (hstmt) != SQL_NO_DATA_FOUND)
{
    // do the processing
}
SQLFreeStmt (hstmt, SQL_CLOSE);
```

Example 8.2 This example shows how to enable the TT_PREFETCH_CLOSE option with JDBC:

```
con = DriverManager.getConnection ("jdbc:timesten:client:" + DSN);
stmt = con.createStatement();

import com.timesten.sql
.....
.....
con.setTtPrefetchClose(true);
rs = stmt.executeQuery("select * from t");
while(rs.next())
{
```

```
// do the processing
}

import com.timesten.sql
....
try {
    ((TimesTenConnection)con).setTtPrefetchClose(true);
}
catch (SQLException) {
    ...
}

rs.close();
con.commit();
```

Use a connection handle when calling SQLTransact

Performance impact: Large

An application can make a call to SQLTransact with either SQL_NULL_HDBC and a valid environment handle:

```
SQLTransact (ValidHENV, SQL_NULL_HDBC, fType)
```

or a valid connection handle:

```
SQLTransact (SQL_NULL_HENV, ValidHDBC, fType).
```

If the intention of the application is simply to commit/rollback on a single connection, it should use a valid connection handle when calling SQLTransact.

SQL tuning

After you have determined the overall logging, locking and I/O strategies, make sure that the individual SQL statements are executed as efficiently as possible.

Tune statements and use indexes

Performance impact: Large

Verify that all statements are executed efficiently. For example, use queries that reference only the rows necessary to produce the required result set. If only `col1` from table `t1` is needed, use the statement:

```
SELECT col1 FROM t1...
```

instead of using:

```
SELECT * FROM t1...
```

[Chapter 9, “The TimesTen Query Optimizer,”](#) describes how to view the plan that TimesTen uses to execute a statement. Alternatively, you can use the `ttIsql showplan` command to view the plan. View the plan for each frequently executed statement in the application. If indexes are not used to evaluate predicates, consider creating new indexes or rewriting the statement or query so that indexes can be used. For example, indexes can only be used to evaluate WHERE clauses when single columns appear on one side of a comparison predicate (equalities and inequalities), or in a BETWEEN predicate.

If this comparison predicate is evaluated often, it would therefore make sense to rewrite

```
WHERE c1+10 < c2+20
```

to

```
WHERE c1 < c2+10
```

and create an index on `c1`.

The presence of indexes does slow down write operations such as [UPDATE](#), [INSERT](#), [DELETE](#) and [CREATE VIEW](#). If an application does few reads but many writes to a table, an index on that table may hurt overall performance rather than help it.

The `FIRST` keyword can be used to operate on a specific number of rows in the SQL statements, `SELECT`, `UPDATE` and `DELETE`. This attribute can improve throughput and response time. Alternatively, if an application plans to fetch at most one row for a query, and a unique index is not being used to fetch the row, the application should set `SQL_MAX_ROW_COUNT` to 1. See the [Oracle TimesTen In-Memory Database API Reference Guide](#).

Occasionally, the system may create a temporary index to speed up query evaluation. If this happens frequently, it is better for the application itself to

create the index. The `CMD_TEMP_INDEXES` column in the `MONITOR` table indicates how often a temporary index was created during query evaluation.

If you have implemented time-based aging for a table or cache group, create an index on the timestamp column for better performance of aging. See [“Time-based aging” on page 133](#).

Select hash or T-tree indexes appropriately

Performance impact: Variable

The TimesTen Data Manager supports both hash and T-tree indexes. Hash indexes are created for tables that declare a `PRIMARY KEY`. T-tree indexes are created with the `CREATE INDEX` statement.

The index structures have different strengths. Hash indexes are faster than T-tree indexes for exact key lookups on all columns of the primary key. However, hash indexes have the following restrictions:

- Hash indexes do not speed up queries on `LESS THAN` or `GREATER THAN` comparisons on indexed columns.
- All columns must appear in the `WHERE` clause.
- Hash indexes can only be created on the columns of the `PRIMARY KEY`; these columns cannot be set to `NULL`, cannot be changed once a record is inserted and only allow unique values of the indexed columns to be inserted.

T-tree indexes can also speed up exact key lookups but are more flexible and can speed up other queries as well. Select a T-tree index if your queries include `LESS THAN` or `GREATER THAN` comparisons. T-tree indexes can also be used to speed up “prefix” queries. A prefix query has equality conditions on all but the last key column that is specified. The last column of a prefix query can have either an equality condition or an inequality condition.

Consider the following table and index definitions:

```
CREATE TABLE T(i1 integer, i2 integer, i3 integer, ...);  
CREATE INDEX IXT on T(i1, i2, i3);
```

The index `IXT` can be used to speed up the following queries:

```
SELECT * FROM T WHERE i1>12;  
SELECT * FROM T WHERE i1=12 and i2=75;  
SELECT * FROM T WHERE i1=12 and i2 BETWEEN 10 and 20;  
SELECT * FROM T WHERE i1=12 and i2=75 and i3>30;
```

The index `IXT` will not be used for queries like:

```
SELECT * FROM T WHERE i2=12;
```

because the prefix property is not satisfied. There is no equality condition for `i1`.

The index `IXT` will be used, but matching will only occur on the first two col-

umns for queries like:

```
SELECT * FROM T WHERE i1=12 and i2<50 and i3=630;
```

T-tree indexes have a dynamic structure that adjusts itself automatically to accommodate changes in table size. A T-tree index can be either unique or non-unique and can be declared over nullable columns. It also allows the indexed column values to be changed once a record is inserted. A T-tree index is likely to be more compact than an equivalent hash index.

Size hash indexes appropriately

Performance impact: Variable

TimesTen uses hash indexes to enforce primary key constraints. The number of buckets used for the hash index is determined by the `PAGES` parameter specified in the `UNIQUE HASH ON` clause of the `CREATE TABLE` statement. The value for `PAGES` should be the expected number of rows in the table divided by 256. A smaller value may result in a greater number of collisions, decreasing performance, while a larger value may provide somewhat increased performance at the cost of extra space used by the index.

If the number of values to be indexed varies dramatically, it is best to err on the side of a large index. If the size of a table cannot be accurately predicted, consider using a T-tree index with `CREATE INDEX`. Also, consider the use of unique indexes when the indexed columns are large CHAR or binary values or when many columns are indexed. Unique indexes may be faster than hash indexes in these cases.

If the performance of record inserts degrades as the size of the table gets larger, it is very likely that you have underestimated the expected size of the table. You can resize the hash index by using the `ALTER TABLE` statement to reset the `PAGES` value in the `UNIQUE HASH ON` clause.

Use foreign key constraint appropriately

Performance impact: Variable

The declaration of a foreign key has no performance impact on `SELECT` queries, but it slows down the `INSERT` and `UPDATE` operations on the table that the foreign key is defined on and the `UPDATE` and `DELETE` operations on the table referenced by the foreign key. The slow down is proportional to the number of foreign keys that either reference or are defined on the table.

Computing exact or estimated statistics

Performance impact: Variable

If statistics are available on the data in the data store, the TimesTen optimizer uses them when preparing a command to determine the optimal path to the data.

If there are no statistics, the optimizer uses generic guesses about the data distribution. For performance reasons, TimesTen does not hold a lock on tables or rows when computing statistics.

If you have examined the plans generated for your statements (see [Chapter 9, “The TimesTen Query Optimizer”](#)) and you think they may not be optimal, consider computing statistics before preparing your statements and re-examining the plans.

If you haven't examined the plans, we generally recommend computing statistics since the information is likely to result in more efficient plans.

There are two built-in procedures for computing statistics: **ttOptUpdateStats** and **ttOptEstimateStats**. **ttOptUpdateStats** looks at every row of the table(s) in question and computes exact statistics. **ttOptEstimateStats** looks at only a sampling of the rows of the table(s) in question and produces estimated statistics. Estimating statistics can be considerably faster but can also result in less accurate statistics. In general, if time is not an issue, it's best to call **ttOptUpdateStats**. But estimation is preferable if overall application performance may be affected. As a rule of thumb, computing statistics with a sample of 10 percent is about ten times faster than computing exact statistics and generally results in the same execution plans. Since computing statistics is a time-consuming operation, it's best to compute statistics once after loading your data store (and before preparing commands), and then periodically only if the composition of your data changes substantially.

Avoid ALTER TABLE

Performance impact: Variable

The ALTER TABLE statement allows applications to add columns to a table and to drop columns from a table. Although the ALTER TABLE statement itself runs very quickly in most cases, the modifications it makes to the table can cause subsequent operations on the table to run more slowly. The actual performance degradation the application experiences varies with the number of times the table has been altered and with the particular operation being performed on the table.

Dropping VARCHAR and VARBINARY columns is slower than dropping columns of other data types since a table scan is required to free the space allocated to the existing VARCHAR and VARBINARY values in the column to be dropped.

Avoid nested queries

Performance impact: Variable

TimesTen supports nested queries with some limitations. However, performance varies and is generally not optimal. See the [Oracle TimesTen In-Memory Database SQL Reference Guide](#) for details on subqueries.

Improving performance of materialized views

Limit number of join rows

Performance impact: Variable

Larger numbers of join rows decrease performance. You can limit the number of join rows by controlling the join condition. For example, use only equality conditions that map one row from one table to one or at most a few rows from the other table.

Use indexes on join columns

Performance impact: Variable

Create indexes on the columns of the detail table that are specified in the `SELECT` statement that creates the join. Also consider creating an index on the materialized view itself. This can improve the performance of keeping the materialized view updated.

If an `UPDATE` or `DELETE` operation on a detail table is often based on a condition on a column, try to create an index on the materialized view on this column if possible.

For example, `CustOrder` is a materialized view of customer orders, based on two tables. The tables are `Customer` and `bookOrder`. The former has two columns (`custNo` and `custName`) and the latter has three columns (`ordNo`, `book`, and `custNo`). If you often update the `bookOrder` table to change a particular order by using the condition `bookOrder.ordNo=const`, then create an index on `CustOrder.ordNo`. On the other hand, if you often update based on the condition `bookOrder.custNo=const`, then create an index on `CustOrder.custNo`.

If you often `UPDATE` using both conditions and cannot afford to create both indexes, you may want to add `bookOrder.rowId` in the view and create an index on it instead. In this case, `TimesTen` updates the view for each detail row update instead of updating all of the rows in the view directly and at the same time. The scan to find the row to be updated is an index scan instead of a row scan, and no join rows need to be generated.

If `ViewUniqueMatchScan` is used in the execution plan, it is a sign that the execution may be slower or require more space than necessary. A `ViewUniqueMatchScan` is used to handle an update or delete that cannot be translated to a direct update or delete of a materialized view, and there is no unique mapping between a join row and the associated row in the materialized view. This can be fixed by selecting a unique key for each detail table that is updated or deleted.

Avoid unnecessary updates

Performance impact: Variable

Try not to update a join column or a “group by” column because this involves deleting the old value and inserting the new value.

Try not to update an expression that references more than one table. This may disallow direct update of the view because TimesTen may perform another join operation to get the new value when one value in this expression is updated.

View maintenance based on an update or delete is more expensive when:

- The view cannot be updated directly. For example, not all columns specified in the detail table UPDATE or DELETE statement are selected in the view, or
- There is not an indication of a one-to-one mapping from the view rows to the join rows.

For example:

```
CREATE MATERIALIZED VIEW v1 AS SELECT x1 FROM t1, t2 WHERE x1=x2;  
DELETE FROM t1 WHERE y1=1;
```

The extra cost comes from the fact that extra processing is needed to ensure that one and only one view row is affected due to a join row.

The problem is resolved if either `x1` is `UNIQUE` or a unique key from `t1` is included in the select list of the view. `ROWID` can always be used as the unique key.

Avoid changes to the inner table of an outer join

Performance impact: Variable

Since outer join maintenance is more expensive when changes happen to an inner table, try to avoid changes to the inner table of an outer join. When possible, perform `INSERT` operations on an inner table before inserting into the associated join rows into an outer table. Likewise, when possible perform `DELETE` operations on the outer table before deleting from the inner table. This avoids having to convert non-matching rows into matching rows or vice versa.

Limit number of columns in a view table

Performance impact: Variable

The number of columns projected in the view `SelectList` can impact performance. As the number of columns in the select list grows, the time to prepare operations on detail tables increases. In addition, the time to execute operations on the view detail tables also increases. Do not select values or expressions that are not needed.

The optimizer considers the use of temporary indexes when preparing operations on detail tables of views. This can significantly slow down prepare time,

depending upon the operation and the view. If prepare time seems slow, consider using `ttOptSetFlag` to turn off temporary tree indexes and temporary hash scans.

Scaling to Multiple CPUs

Run the demo applications as a prototype

Performance impact: Variable

One way to see what sort of scaling you can expect from TimesTen is to run one of the scalable demo applications, such as `tpbmn`, on your system.

The `tpbmn` application has a parameter `-proc` that lets you vary the number of processes that execute TimesTen operations. In addition, it has other parameters that let you specify the transaction mix of READs, WRITEs and INSERTs.

Run `tpbmn -help` to see the full list of options. For example, you can specify whether it uses a hash index or a T-tree index.

By default the demo does one operation per transaction. You can use the `-ops` options to add more operations in a transaction to better model your application. Larger transactions may scale better or worse, depending on the application profile.

Run multi-processor versions of the demo to evaluate how your application can be expected to perform on systems that have multiple CPUs. If the demo scales well, but your application scales poorly, you might try simplifying your application to see where the issue is. Some users “stub out” the TimesTen calls and find they still have bad scaling, and they discover an issue in their application.

You may find that some simulated application data is not being generated properly, so all the operations are accessing the same few rows. That type of localized access will greatly inhibit scalability if the accesses involve changes to the data.

Limit database-intensive connections per CPU

Performance impact: Variable

Check the `LOCK_TIMEOUTS` or `LOCK_GRANTS_WAIT` fields in the `SYS.MONITOR` table. If they have high values, this may indicate undue contention, which can lead to poor scaling.

Because TimesTen is quite CPU-intensive, optimal scaling is achieved by having at most one database-intensive connection per CPU. If you have a 4-CPU system or a 2-CPU system with hyperthreading, then a 4-processor application will run well, but an 8-processor application will not perform well. The contention between the active threads will be too high. The only exception to this rule is when many transactions are committed durably. In this case, the connections are not very CPU-intensive because of the increase in I/O operations to disk, and so the machine can support many more concurrent connections.

Use read operations when available

Performance impact: Variable

Read operations scale better than write operations. Make sure that the read/write balance reflects the real-life workload of your application.

Limit prepares, re-prepares and connects

Performance impact: Variable

Prepares do not scale. Make sure that you pre-prepare commands that are executed more than once. The `CMD_PREPARES` and `CMD_REPREPARES` columns of the `SYS.MONITOR` table indicate how often commands were prepared or automatically re-prepared (due to creation or deletion of indexes). If either has a high value, modify your application to do connection pooling, so that connects and disconnects are rare events.

Connects do not scale. Make sure that you pre-prepare commands that are executed more than once. Look at the `DS_CONNECTS` field in the `SYS.MONITOR` table. If the field has a high value, modify your application to do connection pooling, so that connects and disconnects are rare events.

Limit replication transmitters and receivers and XLA readers

Performance impact: Variable

Replication and XLA operations have significant logging overhead. Replication scales best when there are a limited number of transmitters and/or receivers. Check your replication topology and see if you can simplify it. Generally, XLA scales best when there are a limited number of readers. If your application has numerous readers, see if you can reduce the number.

Monitor XLA and replication to ensure they are reading from the log buffer rather than from the disk. With a lot of concurrent updates, replication may not keep up. Updates are single-threaded at the subscriber. You can achieve better XLA throughput if the frequency of acknowledgements is reduced.

Estimate the number of readers and transmitters required by checking the values in the `LOG_FS_READS` and `LOG_BUFFER_WAITS` columns in the `SYS.MONITOR` table. The system updates this information each time a connection is made or released and each time a transaction is committed or rolled back.

Setting `LogFlushMethod=2` can improve performance of `RETURN TWOSAFE` replication operations and `RETURN RECEIPT` with `DURABLE TRANSMIT` operations.

Allow indexes to be rebuilt in parallel during recovery

Performance impact: Variable

On multi-processor systems, set **RecoveryThreads** to minimum(number of CPUs available, number of indexes) to allow indexes to be rebuilt in parallel if recovery is necessary. If a rebuild is necessary, progress can be viewed in the user log. Setting **RecoveryThreads** to a number larger than the number of CPUs available can cause recovery to take longer than if it were single-threaded.

Use private commands

Performance impact: Variable

On multi-processor systems, if many threads are executing the same commands, then try setting **PrivateCommands=1** to improve throughput or response time. The use of private commands increases the amount of temporary space used.

XLA acknowledgement modes

Prefetch multiple update records

Performance impact: Medium

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 the application can tolerate not acknowledging each message individually.

Acknowledge XLA updates

Performance impact: Medium

To explicitly acknowledge an XLA update, you call `acknowledge` on the update message. Acknowledging a message implicitly acknowledges all previous messages. Typically, you receive and process multiple update messages between acknowledgements. 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.

The TimesTen Query Optimizer

The TimesTen cost-based query optimizer uses information about an application's tables and their available indexes to choose a fast path to the data. Application developers can examine the plan chosen by the optimizer to check that indexes are used appropriately. If necessary, application developers can also modify the optimizer's behavior so that it chooses a different plan.

This chapter includes the following topics:

- [When optimization occurs](#)
- [Viewing a plan](#)
- [Modifying plan generation](#)

When optimization occurs

It is useful to understand when TimesTen performs query optimization, since a single command may be optimized several times.

TimesTen invokes the optimizer whenever a **SELECT**, **UPDATE**, **DELETE**, **INSERT SELECT** or **CREATE MATERIALIZED VIEW** statement is prepared through an ODBC **SQLPrepare** or **SQLExecDirect** function or any of the JDBC execute methods. The resulting plan persists until an *invalidating* event occurs, or the command is dropped by the application. A command is *invalidated* under the following circumstances:

- A table it uses is dropped
- A table it uses is altered
- An index on a table it references is dropped
- An index is created on a table it references
- Statistics are recomputed

An invalid command is usually reprepared automatically just before it is re-executed. This means that the optimizer is invoked again at this time, possibly resulting in a new plan. Thus, a single command may be prepared several times.

Note: When using JDBC, you must manually reprepare commands when a table has been altered.

A command may have to be prepared manually if, for example, the table that the command referenced was dropped and a new table with the same name was created. When you prepare a statement manually, you should commit the prepare statement so it can be shared. If the command is recompiled because it was invalid, and if recompilation involves DDL on one of the referenced tables, then the prepared statement must be committed to release the command lock.

For example, in ODBC a command joining tables T1 and T2 may undergo the following changes:

Action	Description
SQLPrepare	Command is prepared.
SQLExecute	Command is executed.
SQLExecute	Command is re-executed.
.	.
.	.
.	.
Create Index on T1	Command is invalidated.

Action	Description
SQLPrepare	Command is prepared.
SQLExecute	Command is reprepared, then executed.
SQLExecute	Command is re-executed.
.	.
.	.
.	.
ttOptUpdateStats on T1	Command is invalidated (if the invalidate flag is passed to the ttOptUpdateStats procedure).
.	.
.	.
.	.
SQLExecute	Command is reprepared, then executed.
SQLExecute	Command is re-executed.
.	.
.	.
.	.
SQLFreeStmt	Command is dropped.

In JDBC, a command joining tables T1 and T2 may undergo the following changes:

Action	Description
Connection.prepareStatement	Command is prepared.
PreparedStatement.execute	Command is executed.
PreparedStatement.execute	Command is re-executed.
.	.
.	.
.	.
Create Index on T1	Command is invalidated.
PreparedStatement.execute	Command is reprepared, then executed.

Action	Description
Connection.prepareStatement	Command is prepared.
PreparedStatement.execute	Command is re-executed.
.	.
.	.
.	.
ttOptUpdateStats on T1	Command is invalidated (if the invalidate flag is passed to the ttOptUpdateStats procedure).
.	.
.	.
.	.
PreparedStatement.execute	Command is reprepared, then executed.
PreparedStatement.execute	Command is re-executed.
.	.
.	.
.	.
PreparedStatement.close	Command is dropped.

As illustrated, optimization is generally performed at prepare time, but it may also be performed later when indexes are dropped or created, or when statistics are modified. Optimization does not occur if a prepare can use a command in the cache.

If a command was prepared with the `genPlan` flag set, it will be recompiled with the same flag set. Thus, the plan is generated even though the plan for another query was found in the `SYS.PLAN` table.

If an application specifies hints to modify the optimizer's behavior (see [“Modifying plan generation” on page 193](#)), these hints persist until the command is deleted. (For example, when the ODBC `SQLPrepare` function or JDBC `Connection.prepareStatement` method is called again on the same handle or when the `SQLFreeStmt` function or `PreparedStatement.close` method is called.) This means that any intermediate reprepare operations that occur because of invalidations will use those same hints.

Viewing a plan

Several steps are needed to view a plan for a prepared **SELECT**, **UPDATE**, **DELETE**, **INSERT SELECT** or **CREATE VIEW** statement:

- Generating the plan involves instructing TimesTen to store the command's plan in the system **PLAN** table.
- Preparing the statement means calling the ODBC **SQLPrepare** function or JDBC **Connection.prepareStatement** method on the statement.
- Reading the **SYS.PLAN** table is the final step.

After the first two steps have been completed, TimesTen places the command's plan in the **PLAN** table. The stored plan is then updated automatically whenever the command is reprepared. This re-preparation is not likely to occur often. However, if a table in the statement is altered, or if indexes are created or dropped, or the application chooses to invalidate commands when statistics are updated, it may help to read the **PLAN** table again to see if the command's plan has changed.

Generating the plan

Before you can view the plan, you must call the built-in procedure **ttOptSetFlag** with the **genPlan** flag. This call informs TimesTen that all subsequent calls to the ODBC **SQLPrepare** function or JDBC **Connection.prepareStatement** method in the transaction should store the resulting plan in the current **SYS.PLAN** table.

Note: Make sure **AUTOCOMMIT** is not set. If it is, the current transaction completes after the processing of the command and prepares in the next transaction are not affected.

The **SYS.PLAN** table only stores one plan, so each call to the ODBC **SQLPrepare** function or JDBC **Connection.prepareStatement** method overwrites any plan currently stored in the table.

If a command was prepared with the **genPlan** flag set, it will be recompiled with the same flag set. Thus, the plan is generated even though the plan for another query was found in the **SYS.PLAN** table.

For the purposes of experimentation, you can try query and optimizer hints using the **ttIsql** utility. To display optimizer plans, issue the commands:

```
autocommit 0;  
showplan 1;
```

Reading the PLAN table

Once plan generation has been turned on and a command has been prepared, one or more rows in the **SYS.PLAN** table store the plan for the command. The number of rows in the table depends on the complexity of the command. Each

row has seven columns, as described in the chapter “[System and Replication Tables](#)” in the *Oracle TimesTen In-Memory Database API Reference Guide*.

Example 9.1 A
Query and its
plan

Assume you prepare the following query:

```
SELECT COUNT(*)
FROM T1, T2, T3
WHERE T3.B/T1.B > 1
AND T2.B <> 0
AND T1.A = -T2.A
AND T2.A = T3.A
```

The optimizer may generate the five **SYS.PLAN** rows shown in the following table. Each row is one step in the plan and reflects an operation that is performed during query execution.

Step	Level	Operations	Tbl Names	IXName	Pred	Other Pred
1	3	TblLkTtreeScan	T1	IX1		
2	3	TblLkTtreeScan	T2	IX2(D)		T2.B <> 0
3	2	MergeJoin			T1.A = -T2.A	
4	2	TblLkTtreeScan	T3	IX3(D)		
5	1	MergeJoin			T2.A = T3.A	T3.B / T1.B > 1

The remainder of this section provides detailed information about each column in the **SYS.PLAN** table, using [Example 9.1](#) throughout.

PLAN table columns


The [SYS.PLAN](#) table has seven columns. They are described in this section.

Column 1 (Step) Indicates the order of operation. This example uses a table lock T-tree scan. The order is:

1. Table locking T-tree scan of IX1 on table T1.
2. Table locking T-tree scan of IX2 on T2.
3. Merge join of T1 and T2 and so forth.

The steps always start with 1.

Column 2 (Level) Indicates the position of the operation in the join-tree diagram that describes the execution. For the example above, the join tree looks like this:



Column 3 (Operation) Indicates the type of operation being executed. For a description of the values in this field and the type of table scan each represents, see [SYS.PLAN](#) in the chapter “[System and Replication Tables](#)” in the *Oracle TimesTen In-Memory Database API Reference Guide*.

Not all operations the optimizer performs are visible to the user. Only operations significant to performance analysis are shown in the [SYS.PLAN](#) table. TblLk is an optimizer hint that is honored at execution time in serializable or read-committed isolation. Table locks are used during a scan only if row locks are disabled during preparation.

Column 4 (TblNames) Indicates the table that is being scanned. This column is used only when the operation is a scan (one of the first five operations listed above). In all other cases, this column is NULL.

Column 5 (IXName) Indicates the index that is being used. This column is used only when the operation is an index scan using an existing index (using a hash or T-tree scan). In all other cases, this column is NULL. Names of T-tree indexes are followed with “(D)” if the scan is descending (from large to small rather than from small to large).

**Column 6
(Pred)**

Indicates the predicate that participates in the operation, if there is one. Predicates are used only with index scan and MergeJoin operations.

This column may be NULL (no predicate) for a T-tree scan. The optimizer may choose a T-tree scan over a table scan because it has two useful properties in addition to filtering:

- Rows are returned in sorted order (on index key).
- Rows may be returned faster (especially if the table is sparse).

In [Example 9.1](#), the T-tree scans are used for their sorting capability; none of them evaluates a predicate.

The predicate character string is limited to 1,024 characters.

**Column 7
(Other
Pred)**

Indicates any other predicate that is applied while the operation is being executed. These predicates do not participate directly in the scan or join but are evaluated on each row returned by the scan or join.

For example, at step 2 of the plan generated for the example above, a T-tree scan is performed on table T2. When that scan is performed, the predicate `T2.B <> 0` is also evaluated. Similarly, once the final merge-join has been performed, it is then possible to evaluate the predicate `T3.B / T1.B > 1`.

Modifying plan generation

This section explains why you may want to modify execution plans and then describes how to modify them.

Why modify an execution plan?

Applications may want to modify an execution plan for two reasons:

- ***The plan is optimally fast but is ill-suited for the application.*** The optimizer may select the fastest execution path, but this path may not be desirable from the application's point of view. For example, if the optimizer chooses to use certain indexes, these choices may prevent other operations—such as certain update or delete operations—from occurring simultaneously on the indexed tables. In this case, an application can prevent the use of those indexes.

The plan chosen by the optimizer may also consume more memory than is available or than the application wants to allocate. For example, this may happen if the plan stores intermediate results or requires the creation of temporary indexes.

- ***The plan is not optimally fast.*** The query optimizer chooses the plan that it estimates will execute fastest based on its knowledge of the tables' contents, available indexes, statistics and the relative costs of various internal operations. The optimizer often has to make estimates or generalizations when evaluating this information, so there can be instances where it does not choose the fastest plan. In this case, an application can adjust the optimizer's behavior to try to produce a better plan.

When to modify an execution plan

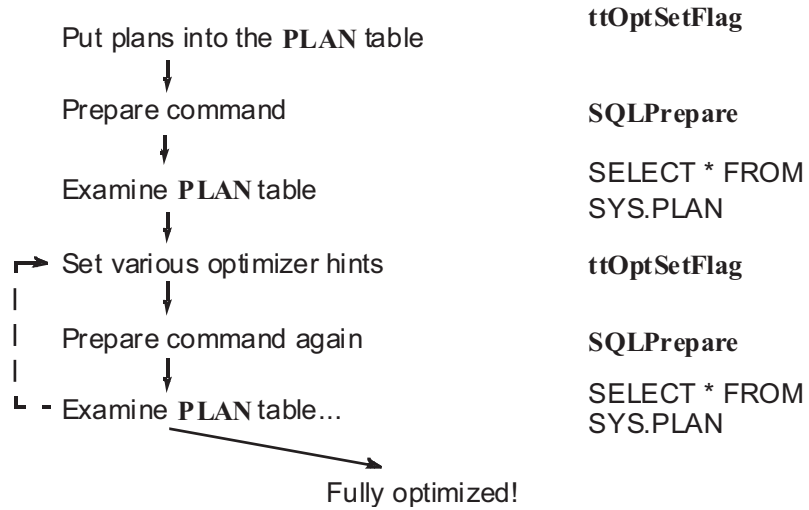
Applications can modify an execution plan by giving hints to the optimizer. Hints are specified by calls to one of the TimesTen optimizer built-in procedures (see [“How to modify execution plan generation” on page 197](#)) and are in effect for all calls to the ODBC **SQLPrepare** function or JDBC **PreparedStatement** objects in the transaction.

Note: Make sure AUTOCOMMIT is not set. If it is, the current transaction completes after processing the **ttOptSetFlag** procedure and prepares in the next transaction are not affected.

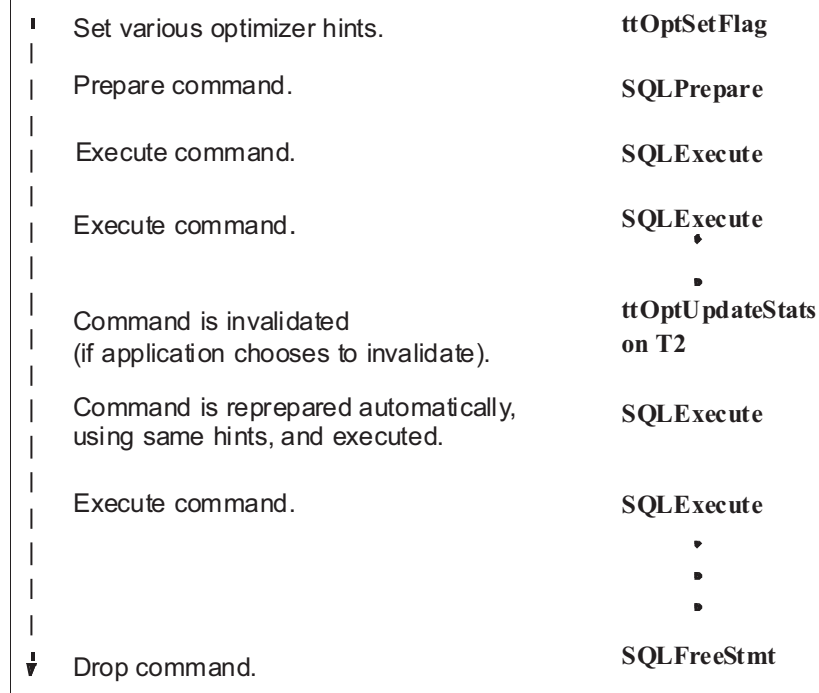
If a command is prepared with certain hints in effect, those hints continue to apply if the command is reprepared automatically, even when this happens outside the initial prepare's transaction. This can happen when a table is altered, or an index is dropped or created, or when statistics are modified, as described in [“When optimization occurs” on page 186](#).

If a command is prepared without hints, subsequent hints will not affect the command if it is reprepared automatically. An application must call the ODBC **SQLPrepare** function or JDBC **Connection.prepareStatement** method a second time so that hints have an effect.

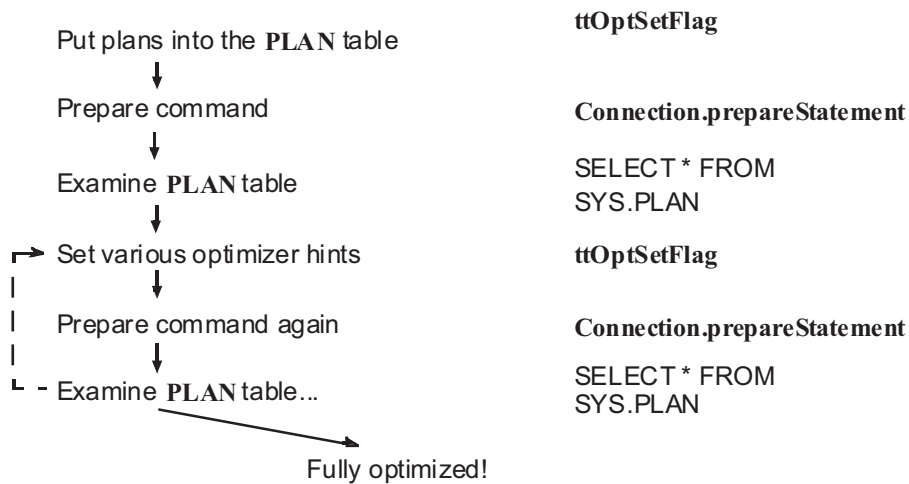
Example 9.2 When using ODBC, a developer tuning a join on T1 and T2 might go through the following steps:



During execution, the application may then go through the following steps (with no user intervention):



Example 9.3 When using JDBC, a developer tuning a join on T1 and T2 might go through the following steps:



During execution, the application may then go through the following steps (with no user intervention):

Set various optimizer hints.	ttOptSetFlag
Prepare command.	Connection.prepareStatement
Execute command.	Statement*.execute*
Execute command.	Statement*.execute* • •
Command is invalidated (if application chooses to invalidate).	ttOptUpdateStats on T2
Command is re-prepared automatically, using same hints, and executed.	Statement*.execute*
Execute command.	Statement*.execute* • • •
Drop command.	PreparedStatement.close

How to modify execution plan generation

To change the query optimizer behavior, an application calls one of the following built-in procedures using the ODBC procedure call interface:

- [ttOptClearStats](#)
- [ttOptEstimateStats](#)
- [ttOptSetColIntvlStats](#)
- [ttOptSetFlag](#)
- [ttOptSetOrder](#)
- [ttOptSetTblStats](#)
- [ttOptUpdateStats](#)
- [ttOptUseIndex](#)

The procedure [ttOptSetFlag](#) sets certain optimizer parameters. [ttOptSetOrder](#) allows an application to specify the table join order. The procedure [ttOptUseIndex](#) allows an application to disable the use of certain indexes. The remaining procedures manipulate statistics the TimesTen Data Manager maintains on the application's data that are used by the query optimizer to estimate costs of various operations. See [Chapter 2, "Built-In Procedures"](#) in the *Oracle TimesTen In-Memory Database API Reference Guide*.

[Example 9.4](#) shows an ODBC example of how to use [ttOptSetFlag](#). [Example 9.5](#) shows a JDBC example.

Example 9.4 The ODBC example below illustrates the use of [ttOptSetFlag](#) to prevent the optimizer from choosing a merge join.

```
import java.sql.*;
class Example
{
    public void myMethod() {
        CallableStatement cStmt;
        PreparedStatement pStmt;
        . . . . .
        try {
            . . . . .
            // Prevent the optimizer from choosing Merge Join
            cStmt = con.prepareCall("{
                CALL ttOptSetFlag('MergeJoin', 0)}");
            cStmt.execute();
        }
    }
}
```

```

        // Next prepared query
        pstmt=con.prepareStatement(
        "SELECT * FROM Tbl1, Tbl2 WHERE Tbl1.ssn=Tbl2.ssn");
        . . . . .
        catch (SQLException ex) {
            ex.printStackTrace();
        }
    }
    . . . . .
}

```

Example 9.5 The JDBC example below illustrates the use of **ttOptSetFlag** to prevent the optimizer from choosing a merge join.

```

#include <sql.h>
SQLRETURN rc;
SQLHSTMT hstmt; fetchStmt;
....
rc = SQLExecDirect (hstmt, (SQLCHAR *)
    "{CALL ttOptSetFlag (MergeJoin, 0)}",
    SQL_NTS)
/* check return value */
...
rc = SQLPrepare (fetchStmt, ...)
/* check return value */
...

```

You can also experiment with optimizer settings using the **ttIsql** utility. The commands that start with **try** control the optimizer hints. To view the current optimizer hint settings, use the **optprofile** command.

Glossary

.odbc.ini file

See ODBC.INI file.

ACID transaction semantics

An acronym referring to the four fundamental properties of a transaction: atomicity, consistency, isolation and durability.

atomicity

A property of a transaction whereby either all or none of the operations of a transaction are applied to the data store.

backup instance

A set of files containing backup information for a given data store, residing at a given [“backup path”](#). See also [“backup point”](#), [“full backup”](#) and [“incremental backup”](#).

backup path

The location of a data store, specified by a directory name and an optional basename.

backup point

The time at which a backup begins. See also [“backup path”](#), [“full backup”](#) and [“incremental backup”](#).

cache group

A set of cached tables related through foreign keys.

cache instance

A set of rows related through foreign keys. Each cache instance contains exactly one row from the root table of a cache group and zero or more rows from the other tables in the cache group.

client/server

An approach to application design and development in which application processing is divided between components running on an end user's machine (the client) and a network server. Generally, user interface elements are implemented in the client component, while the server controls database access.

client data source name

See [“data source name, client”](#).

concurrency

The ability to have multiple transactions access and manipulate the data store at the same time.

connection

A data path between an application and a particular ODBC data source.

connection attribute

A character string that defines a connection parameter to be used when connecting to an ODBC data source. Connection attributes have the form *name=value*, where *name* is the name of the parameter and *value* is the parameter value. See also connection string.

connection request

A message sent by an application through an ODBC driver to an ODBC data source to request a connection to that data source.

connection string

A character string that defines the connection parameters to be used when connecting to an ODBC data source. A connection string is expressed as one or more connection attributes separated by semicolons.

consistency

A property of transactions whereby each transaction transforms the database from one consistent state to another.

cursor

A control structure used by an application to iterate through the results of an SQL query.

data source definition

A named collection of connection attributes that defines the connection parameters to be used when connecting to an ODBC data source. See also “[data source name](#)”.

data source name

A logical name by which an end user or application refers to an ODBC data source definition. Sometimes incorrectly used to mean “data source definition.” See also “[data source definition](#)”, ODBC.INI file.

data source name, client

A data source name defined on a TimesTen client machine that refers to a TimesTen server data source name on a server machine.

data source name, server

A system data source name (system DSN) defined on a server machine. Server Data Source Names become available to all TimesTen clients on a network when the TimesTen Server is running.

data source name, system

A data source name that is accessible by all users of a particular machine.

data source name, user

A data source name that is accessible only by the user who created the data source name.

driver

See [“ODBC driver”](#).

DSN

See [“data source name”](#).

DSN, client

See [“data source name, client”](#).

DSN, server

See [“data source name, server”](#).

DSN, system

See [“data source name, system”](#).

DSN, user

See [“data source name, user”](#).

durability

A property of transactions whereby the effects of a committed transaction survive system failures.

environment variable

A *name, value* pair maintained by the operating system that can be used to pass configuration parameters to an application.

event

An activity or occurrence that can be tracked by a logging mechanism in an application, service or operating system. See also [“logging”](#), [“protocol message logging”](#) and [“event viewer”](#).

event viewer

On Windows, a utility program used to view the contents of the operating system event log.

full backup

A data store backup procedure in which a complete copy of a data store is created. Typically, the first backup of a data store must be a full backup. See also [“incremental backup”](#).

host

A computer. Typically used to refer to a computer on a network that provides services to other computers on the network.

host name

A character string name that uniquely identifies a particular computer on a network. Examples: athena, thames.mycompany.com. See also [“host”](#).

in-line column

A column whose values are physically stored together with the other column values of a row.

incremental backup

A data store backup procedure in which an existing backup is augmented with all the log records created since its last full or incremental backup. See also [“backup instance”](#) and [“full backup”](#).

initialization file

See ODBC.INI file.

IP address

A numeric address that uniquely identifies a computer on a network and consists of four numbers separated by dots. Abbreviation for Internet Protocol address. Example: 123.61.129.91.

IPC

Inter Process Communication.

isolation

A property of transactions whereby each transaction runs as if it were the only transaction in the system.

listener thread

A thread that runs on the TimesTen Server that receives and processes connection requests from TimesTen Clients.

logging

The process by which an application, service or operating system records specific events that occur during processing.

multithreading

A programming paradigm in which a process contains multiple threads of control.

network address

A host name, or IP address that uniquely identifies a particular computer on a network. Examples: 123.61.129.91, athena, thams.mycompany.com.

ODBC

See [“Open Database Connectivity \(ODBC\)”](#).

ODBC Administrator

A utility program used on Windows to create, configure and delete data source definitions.

ODBC data source

See [“data source name”](#) (DSN).

ODBC data source name

See [“data source name”](#) (DSN).

ODBC driver

A library that implements the function calls defined in the ODBC API and enables applications to interact with ODBC data sources.

ODBC Driver Manager

A library that acts as an intermediary between an ODBC application and one or more ODBC drivers.

ODBC initialization file

See [“.odbc.ini file”](#).

Open Database Connectivity (ODBC)

A database-independent application programming interface that enables applications to access data stored in heterogeneous relational and non-relational databases. Based on the Call-Level Interface (CLI) specification developed by X/Open's SQL Access Group and first popularized by Microsoft on the Windows platform.

out-of-line column

A column whose values are physically stored separately from the other column values of a row.

phantom

A row that appears during one read but not during another read within the same transaction, due to the actions of other concurrently executing transactions.

ping

A utility that tests the connection between two computers on a network by sending a message from one computer to the other and measuring how long it takes for the receiving system to confirm that the message was received. Typically packaged with network software.

port number

See [“TCP/IP port number”](#).

procedure

See [“stored procedure”](#).

process

An instance of a program in execution.

propagate

When using Cache Connect to Oracle, to send table or row modifications from a Cache Connect to Oracle data store to an Oracle database. Compare with [“replicate”](#).

protocol message logging

The process that the TimesTen Server uses to record each message it receives through the TimesTen network protocol.

replicate

The sending of table or row modifications from one data store to another. Compare with [“propagate”](#).

result set

A collection of zero or more rows of data that represent the result of an SQL query.

rollback

To undo the actions of a transaction, thereby returning all items modified by the transaction to their original state.

row buffering

A performance enhancement used by the TimesTen Client in which the client receives multiple result rows of an SQL query in each message from the TimesTen Server to reduce network communication.

RPC

Remote Procedure Call.

scalability

The degree to which a system or application can handle increasing demands on system resources without significant performance degradation.

server data source name

See [“data source name, server”](#).

server DSN

See [“data source name, server”](#).

system DSN

See [“data source name, system”](#).

shorthand name

A logical name used to refer to a particular TimesTen Server. Shorthand names relieve the end user of having to enter a host name and port number to connect to a TimesTen Server.

SMP

Symmetric multi-processing. A hardware configuration in which two or more similar processors are connected via a high-bandwidth link and managed by one operating system, where each processor has equal access to I/O devices.

SNMP

Simple Network Management Protocol. Used to manage nodes on a network.

SQL

Structured Query Language.

stack overflow condition

An error condition in which the stack usage of a thread or process exceeds the amount of space allocated for the stack.

stored procedure

An executable object or named entity stored in a data store that can be invoked with input and output parameters and which can return result sets similar to those returned by an SQL query.

system account

A special account on Windows used by the operating system and certain operating system services. The TimesTen Service and the TimesTen Server run under the system account.

system DSN

See [“data source name, system”](#).

T-tree

An index structure similar in functionality to a B-tree but optimized for in-memory data management.

TCP/IP

The communications protocol used by computers on the Internet. Abbreviation for Transport Control Protocol/Internet Protocol.

TCP/IP port number

A number used by TCP/IP that identifies the end point for a connection to a host that supports multiple simultaneous connections.

telnet

A utility program and protocol that enables a user on one computer to open a virtual terminal, log in to a remote host and interact as a terminal user of that host.

thread

An independent sequence of execution of program code inside a process. See also “[process](#)”.

thread-safe ODBC driver

An ODBC driver that supports multithreaded servers and clients. The TimesTen data manager driver and the TimesTen Client driver are thread-safe.

timeout error

An error condition indicating that the requested operation did not complete within the given amount of time. See also “[timeout interval](#)”.

timeout interval

A configuration parameter that specifies the maximum amount of time that an operation should take to complete. See also “[timeout error](#)”.

TimesTen Client

(1) An ODBC driver that enables end users to access data sources through a TimesTen Server. (2) A computer on which the TimesTen Client software has been installed. Using the TimesTen Client driver, an end user or application can access any data source managed by an available TimesTen Server.

TimesTen Client/Server network protocol

The protocol used by TimesTen Clients and TimesTen Servers to exchange data over a standard TCP/IP network connection.

TimesTen Data Server

(1) An application program that makes TimesTen data sources available to the TimesTen Clients on a network. (2) A computer on which the TimesTen Data Server software is running.

TimesTen Server address

The host name or IP address used during installation of the TimesTen Server to identify the computer on which the software is being installed.

transaction

An operation or set of operations performed against data in a data store. The operations defined in a transaction must be completed as a whole; if any part of the transaction fails, the entire transaction fails. See also “[ACID transaction semantics](#)”.

UCS-4

A fixed-width, 32-bit Unicode character set. Each character occupies 32 bits of storage. The UCS-2 characters are the first 65,536 code points in this standard, so it can be viewed as a 32-bit extension of UCS-2.

UTF-16

An encoding scheme defined by the ISO/IEC 10646 standard in which each Unicode character is represented by either a two-byte integer or a pair of two-byte integers. Characters from European scripts and most Asian scripts are represented in two bytes. Surrogate pairs are represented in four bytes. Surrogate pairs represent characters such as infrequently used Asian characters that were not included in the original range of two-byte characters.

user account

The combination of a user name, password and access permissions that gives an individual user access to an operating system.

user data source name

See “[data source name, user](#)”.

user DSN

See “[data source name, user](#)”.

User Manager

A Windows utility program used to create user accounts and assign access rights and group membership.

Windows sockets (Winsock)

An API that defines a standard binary interface for TCP/IP transports on Windows platforms. This API adds Windows-specific extensions to the Berkeley Sockets interface originally defined in Berkeley UNIX.

Index

Symbols

.odbc.ini 35

A

accent-insensitive linguistic sort 80
acknowledging updates 184
active standby pair
 aging 135
adding columns 127
adding rows 127
adding rows to a table 32
aging
 active standby pair 135
 attributes 131
 foreign keys 134
 LRU 131
 ON DELETE CASCADE 134
 parent and child tables 134
 performance 133
 replication 135
 restrictions 131
 tables 131
 time-based 131, 133
 usage-based 131
ALTER 177
ALTER SESSION SQL statement 83
ALTER TABLE
 adding and removing columns 127
 and performance 177
application failure
 use of logs and locks 149
ASCIISTR SQL function 82
asynchronous checkpoints
 See fuzzy checkpoints
attributes
 and data source name 25
 PermWarnThreshold 30
 setting for UNIX 23
 specifying 18
 TempWarnThreshold 30
 viewing and changing by ttIsql 112
AUTO_ACKNOWLEDGE mode 184
autocommit ttIsql command 107
automatic

checkpointing 162
index creation 129

B

backing up a data store 32
batch mode
 ttIsql 97
byte semantics 77

C

Cache Connect attributes 20, 21
cache groups
 time-based aging 133
cachegroups ttIsql command 109
case-insensitive linguistic sort 80
changing shared memory segment size 71
character semantics 77
character set
 and replication 77
 choosing 76
character sets
 working with in ttIsql 106
checkpoints
 automatic 126, 162
 fuzzy 161
 influences on duration 160
 static 161
CHR SQL function 83
clearing command history
 ttIsql 105
Client connection attributes
 described 43, 44
Client DSN
 creating 51
 creating on UNIX 54
 creating on Windows 46
 Data Source Setup dialog 47
 name 47
Client/Server
 changing shared memory segment size 71
 configuring 41
 managing shared memory segment size 70
Client/Server communication
 shared memory 69

- TCP/IP 69
- client/server communication
 - overview 40
 - TCP/IP 40
 - UNIX socket 40
- CLIENT_ACKNOWLEDGE mode 184
- close tTsql command 118
- coexistence of different locking levels 159
- column values
 - default 128
- columns
 - in-line 127
 - nullability 127
- command history
 - tTsql 104
- commit tTsql command 107
- commitdurable tTsql command 107
- components of data store 126
- concurrency
 - and logging 155
 - types of isolation 157
- configuring
 - Client/Server 41
- connection character set 48, 78
- Connection.commit method 149
- Connection.prepareStatement method 189
 - and execution plan generation 194
- Connection.rollback method 149
- Connection.setAutoCommit method 149
- ConnectionCharacterSet general connection attribute 78
- connections
 - performance overhead 167
 - testing on Windows 50
- controlling logging 155
- controlling Server log messages 71
- controlling the TimesTen Server daemon 67
- controlling web server options 72
- copying a data store 32
- CREATE INDEX statement 144
- creating a Client DSN
 - on UNIX 54
 - on Windows 46
- creating a logical server name on UNIX 52
- creating a Server
 - name 45
- creating a Server DSN 43
- creating indexes
 - example 144

- explicit creation 129
 - how to do 144
- creating tables
 - example 130, 136
- Custom setup, Windows 11

D

- daemon
 - control operations 66
 - informational messages on Windows 65
 - overview 59
 - starting and stopping on UNIX 60
 - starting and stopping on Windows 60
- daemon startup script 60, 61
- data
 - permanent 29
 - temporary 29
- Data Manager Service 60
- data source
 - specification 35
 - UNIX configuration files 37
- data source names, *See* DSN
- data store
 - accessing on a local machine 40
 - accessing on a remote machine 40
 - backing up 32
 - changing size of 29
 - components 126
 - copying 32
 - definition 9
 - getting information with tTsql 109
 - migrating 32
 - path names, environment variables in 23
 - permanent 126
 - prefix name 23
 - restoring 32
 - setting attributes for UNIX 23
 - sizing 166
 - TCP/IP client/server access 40
 - temporary 24, 127, 167
 - UID connection attribute 126
 - UNIX socket client/server access 40
 - user names 126
 - users and owners 126
- data store-level locking 159
- database character set
 - application 77
 - client operating system 77
 - languages 76

- performance 77
- DatabaseCharacterSet data store attribute 76
- default column values 128
- deleting
 - rows 145
 - Windows server name 48
- describe ttIsql command 109, 118
- detail table 135
- driver
 - JDBC 17
- driver manager
 - JDBC 12
 - linking with 10
 - ODBC 10
- DRIVER parameter 22
- DriverManager.getConnection 44
- DriverManager.getConnection method
 - methods 44
- dropping
 - indexes 144
 - tables, example 130, 140
- DSN
 - .odbc.ini file 35
 - Client 14
 - connection attributes, Data Manager 15
 - Data Manager 14
 - description of types 15
 - example, Windows
 - examples 24
 - finding in order of precedence 52
 - maximum length 13
 - naming rules 13
 - setting attributes 25
 - system 13
 - user 13
- dssize ttIsql command 109
- DUPS_OK_ACKNOWLEDGE mode 184
- durability
 - overview 151

E

- editline for ttIsql 102
- environment variables
 - in data store path names 23
 - TTISQL 122
- estimating
 - data store size 166
 - index size 144
 - table size 130

- examples

- creating indexes 144
 - creating tables 130, 136
 - dropping tables 130, 140
 - PLAN rows 190
- exec ttIsql command 118
- execandfetch ttIsql command 118
- execution plan
 - generating 186, 194
 - modifying 193, 197
 - viewing 189
- execution plan generation
 - and Connection.prepareStatement 194

F

- fetchall ttIsql command 118
- fetchone ttIsql command 118
- files
 - .ttconnect.ini 57
 - odbc.ini 35, 37
 - ttendaemon.options 61, 66, 70
- first connection attributes 19
- foreign key constraint
 - and performance 176
- fragmentation
 - block-level 166
- free ttIsql command 118
- fuzzy checkpoints, definition 161

G

- general connection attributes 19

H

- hash indexes
 - definition 143
 - sizing 176
 - vs. T-tree indexes 175
 - when used 143

I

- index
 - aging performance 133
- index size
 - estimating 144
 - ttSize utility 144
- indexes
 - and performance 174
 - automatic creation 129

- creating 129, 144
- dropping 144
- overview 142
- owner 144
- referencing 144
- See also* T-tree indexes
- See also* hash indexes
- unique 143
- informational messages
 - modifying 64
- in-line columns 127
- inserting rows 145
- INSTR SQL function 82
- INSTR4 SQL function 82
- INSTRB SQL function 82
- interactive mode
 - ttIsql 97
- invalidating commands 186
- isolation modes
 - described 157
 - SERIALIZABLE 157
- isolation ttIsql command 107
- IXNAME column in PLAN table 191

J

- JDBC
 - driver 12
 - driver manager 12
- JDBC tracing 169
- join
 - columns, view performance 178
 - rows, view performance 178

L

- least recently used aging 131
- length semantics 77
- LENGTH SQL function 82
- LENGTH4 SQL function 82
- LENGTHB SQL function 82
- LEVEL column in PLAN table 191
- linguistic index 81
- linguistic sort
 - monolingual 79
 - multilingual 80
- linguistic sorts
 - accent-insensitive 80
 - case-insensitive 80
- linking applications

- Client/Server 39
 - direct 10
 - UNIX 10
 - Windows 10
 - with driver manager 10
- loading a data store to memory 31
- locking
 - See also* locks
 - See also* ttLockLevel
- locks
 - coexistence of different levels 159
 - overview table 148
- log files
 - names 156
- logging
 - attribute for concurrent connections 155
 - how to control 155
 - overview table 148
 - temporary data store 127
- logical server
 - Network Address 41, 45, 52
- logical server name
 - creating on UNIX 52
- LOWER and UPPER SQL functions 82
- LRU aging 131
- LRU aging attributes 131

M

- maintenance options
 - performance impact 168
- managing shared memory segment size 70
- materialized views, *see* views
- maximum name length 128
- memory
 - loading a data store into 31
 - partitions 29
 - permanent 29
 - policy for loading a data store 31
 - temporary 29
- messages
 - informational 43
 - server log 43
- metadata, TimesTen 129
- methods
 - Connection.commit 149
 - Connection.prepareStatement 189
 - Connection.rollback 149
 - Connection.setAutoCommit 149
- migrating data stores 32

- modifying execution plan
 - overview 193
 - procedure overview 197
- modifying informational messages 64
- modifying Server daemon options 67
- modifying web server options 72
- monitor ttlsq command 109

N

- names
 - log files 156
 - maximum 128
- naming a Client DSN 47
- NCHR SQL function 83
- nested subqueries
 - and performance 177
- Network Address for logical server
 - ttLocalHost 41
 - ttShmHost 40
 - UNIX 52
 - Windows 45
- NLS connection attributes 20
- NLS_LENGTH_SEMANTICS general connection attribute 77, 78
- NLS_SORT connection attribute 81
- NLSSORT SQL function 81
- non-materialized view
 - creating 141
- nonmaterialized view
 - description 140
- non-materialized views 141
 - dropping 142
 - restrictions 142
 - SELECT query 141
- not inline columns 127
- NULL values, sorting 143
- nullable columns
 - definition 127
 - primary key not nullable 129

O

- ODBC
 - tracing and performance 169
 - UNIX driver 22
- ODBC functions
 - and the JDBC driver 12
- odbc.ini 35
 - entry example 37

- format of 37
- online help 100
- OPERATION column in PLAN table 191
- optimizer
 - application hints 188
 - example scenario 186, 187
 - generating plan 189
 - invalid statistics 186
 - modifying execution plan 193
 - PLAN row example 190
 - reading plan 189
 - viewing plans 189

- optprofile ttlsq command 113
- OS paging 168
- OTHERPRED column in PLAN table 192
- outer join
 - materialized view performance 179
- out-of-memory warnings 30
- owners
 - of indexes 144

P

- performance
 - and altered tables 177
 - and autocommit 149
 - and foreign key constraints 176
 - and temporary data stores 166
 - application tuning
 - connection overhead 167
 - maintenance options 168
 - ODBC tracing 169
 - automatic index creation 129
 - Client 170
 - data store tuning
 - driver usage 168
 - specifying size 166
 - temporary vs. permanent data store 167
 - join columns in materialized views 178
 - join rows in materialized views 178
 - materialized views 179
 - SQL tuning
 - indexes 174
 - tuning 165
 - working locally 170
 - permanent data partition 29
 - permanent data store 126
 - automatic checkpointing 126
 - PermWarnThreshold attribute 30
 - plan

- See* execution plan 191
- PLAN rows 190
- PLAN table
 - columns 191
 - IXNAME column 191
 - LEVEL column 191
 - OPERATION column 191
 - OTHERPRED column 192
 - PRED column 192
 - STEP column 191
 - TBLNAME column 191
- PRED column in PLAN table 192
 - limit on length 192
- prefetch multiple update records 184
- prepare ttIsql command 118
- PreparedStatement objects 193
- preparing statements 189
- primary keys 129
 - nullability 129
 - See Also* unique indexes

Q

- query optimizer
 - See* optimizer
- query optimizer plans 186
 - viewing with ttIsql 113

R

- RAM policy
 - defined 31
- referencing indexes 144
- remote data store
 - accessing on UNIX 52, 54, 56
 - accessing on Windows 49
- removing
 - columns 127
 - rows 127
- replication
 - aging 135
 - and TimesTen daemon 59
 - character sets 77
 - temporary data partition 30
 - TTREP system tables 128
- restoring a data store 32
- restrictions on table names 128
- rollback
 - logs and locks 148
- rollback ttIsql command 107

- row-level locking 159
- rows
 - deleting 145
 - in-line and out-of-line portions 128
 - inserting 145
 - understanding 144
- RTRIM SQL function 82

S

- saving command history
 - ttIsql 105
- SERIALIZABLE isolation mode 157
- serializable transactions
 - performance 171
- Server
 - controlling the daemon 67
 - creating a DSN 43
 - described 42
 - modifying daemon 67
 - name 47, 48
 - Server List dialog 45
 - Server Name Setup dialog 45
 - starting and stopping 42
- Server log messages 71
- server name
 - creating on Windows 45
- serverShmIpc 70
- serverShmSize 71
- Service
 - See* daemon
- setjoinorder ttIsql command 113
- setting data store attributes
 - UNIX 23
- setting the timeout interval on Windows 48
- setuseindex ttIsql command 113
- shared data store
 - durable and nondurable connections 155
- shared memory
 - Client/Server IPC 69
 - Client/Server, changing size 71
 - Client/Server, managing size 70
- shared memory IPC-enabled server 70
- showplan command 113, 120
- size
 - data store 29
- sizing
 - data stores 166
 - hash indexes 176
 - sorting NULL values 143

- specifying data store size 166
- SQL
 - tuning and performance 174
- SQLPrepare 189
 - and execution plan generation 194
- sqlquerytimeout ttIsql command 107
- SQLTransact
 - undoing effects 149
- starting and stopping the Server 42
- starting and stopping the TimesTen Data Manager 60
- starting the daemon
 - on UNIX 60
 - on Windows 60
- starting the Data Manager Service 60
- static checkpoints, definition 161
- statistics
 - computing 176
 - recomputation 186
 - ttOptEstimateStats 177
 - ttOptUpdateStats 177
- STEP column in PLAN table 191
- stopping the daemon
 - on UNIX 60
 - on Windows 60
- stopping the Data Manager Service 60
- subdaemons
 - minimum required 66
 - setting allowable number 66
 - specifying allowable range 66
- SUBSTR SQL function 82
- SUBSTR4 SQL function 82
- SUBSTRB SQL function 82
- synchronous checkpoints
 - See* static checkpoints
- SYS owner of tables 128
- syslog 65
- System DSN 17, 49, 50
- system failure
 - resulting logs and locks 149
- system tables
 - indexes on 144
 - overview 129

T

- table size
 - estimating 130
 - ttSize utility 130
- tables
 - adding rows 32
 - creating, example 130, 136
 - deleting rows 145
 - dropping, example 130, 140
 - format 127
 - in-line vs. not inline columns 127
 - inserting rows 145
 - modifying format 127
 - name length 128
 - names 128
 - names, restrictions 128
 - nullable columns 127
 - owners 128
 - SYS owner 128
 - understanding rows 144
 - unique indexes 129
- TBLNAME column in PLAN table 191
- TCP/IP client/server communication 40
- Temporary attribute
 - performance 166, 167
- temporary data partition
 - and replication 30
- temporary data store
 - logging 127
- TempWarnThreshold attribute 30
- testing connections on Windows 50
- thread programming
 - and TimesTen 16
- time-based aging 133
- timeout interval
 - setting on Windows 48
- TimesTen
 - ODBC driver 10
 - thread programming 16
- TimesTen Server daemon 67
- TIMESTEN8 character set 84
- timestend 59
- timing ODBC function calls 117
- transaction commit
 - logs and locks 148
- transaction management
 - isolation levels 147
 - locking 148
 - logging 148
 - semantics 148
- transaction rollback
 - logs and locks 148
- tryhash ttIsql command 113
- trymergejoin ttIsql command 113

- trynestedloopjoin ttIsqL command 113
- tryrowid ttIsqL command 113
- tryserial ttIsqL command 113
- trybllocks ttIsqL command 113
- trytmphash ttIsqL command 113
- trytmptable ttIsqL command 113
- trytmptree ttIsqL command 113
- trytree ttIsqL command 113
- TT_PREFETCH_CLOSE connection option 171
- ttAgingLRUConfig built-in procedure 131
- ttconnect.ini file 57
- ttendaemon.options file 61, 65, 66, 70
- ttIsqL 100
 - built-in command usage 107, 109, 110, 116, 117, 118
 - clearing command history 105
 - command history 104
 - customizing command prompt 99
 - deleting rows 145
 - displaying data store information 109
 - editline feature 102
 - modes, interactive and batch 97
 - online help 100
 - saving command history 105
 - timing ODBC function calls 117
 - using 95
 - view and set data store attributes 112
 - viewing optimizer plan 113
 - working with character sets 106
 - working with parameterized SQL statements 118
 - working with prepared SQL statements 118
 - working with transactions 107
- TTISQL environment variable 122
- ttLocalHost
 - logical server address 41
- ttOptEstimateStats
 - statistics computing 177
- T-tree indexes
 - vs. hash indexes 175
- TTREP system tables 128
- ttShmHost 58
 - logical server address 40
- ttWarnOnLowMemory procedure 30

U

- unique indexes 129, 143

- See also* primary key
- UNISTR SQL function 82
- UNIX
 - setting attributes 23
- UNIX configuration file
 - odbc.ini 37
- UNIX configuration files
 - .odbc.ini, format of 37
- UNIX socket client/server communication 40
- unsetjoinorder ttIsqL command 113
- unsetuseindex ttIsqL command 113
- updates
 - materialized view performance 179
- usage-based aging 131
- using shared memory for Client/Server IPC 69

V

- view
 - creating 141
 - description 140
- viewing query optimizer plans 189
- views
 - creating 136
 - dropping 142
 - non-materialized 141
 - performance 139
 - restrictions 142
 - restrictions on detail tables 138
 - restrictions on view tables 138
 - SELECT queries in 137
 - SELECT query for non-materialized 141
 - understanding 135

W

- web server options
 - modifying 72
- working with
 - parameterized SQL statements 118
 - prepared SQL statements 118

X

- XLA bookmark, deleting 124
- XLA updates
 - acknowledging 184
- xldeletebookmark ttIsqL command 124