

Oracle® TimesTen In-Memory Database

Monitoring and Troubleshooting Guide



Release 22.1
F35395-02
August 2025



Copyright © 2011, 2025, Oracle and/or its affiliates.

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

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

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

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

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

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

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

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

Contents

What's New

New Features in Release 22.1.1.1.0

i

1 The TimesTen Prometheus Exporter

About Prometheus	1
About the TimesTen Exporter	2
About Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication	3
Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication	5
About Configuring the TimesTen Exporter and Prometheus with No Authentication	8
Configuring the TimesTen Exporter and Prometheus with No Authentication	9
The Metrics Supported by the TimesTen Exporter	11
Metrics timesten_aging_lru_commit_total to timesten_cg_autorefresh_row_update_total	12
Metrics timesten_cg_awt_oracle_call_total to timesten_cg_flushed_row_total	13
Metrics timesten_cg_parawt_auto_apply_applctn_seqnum to timesten_cg_parawt_auto_apply_partitions	15
Metrics timesten_cg_parawt_auto_apply_plsqlmode_batches to timesten_cg_parawt_auto_apply_validbatch_waits	16
Metrics timesten_cg_parawt_auto_rx_batches to timesten_cg_parawt_auto_sorter_txns_waits	17
Metrics timesten_cg_parawt_auto_tx_applier_waits to timesten_ckpt_write_total	19
Metrics timesten_collections_timedout to timesten_cs_server_transmitted_byte_total	20
Metrics timesten_cursor_close_total to timesten_index_range_rebuild_insert_total	21
Metrics timesten_index_range_row_fetch_repl_total to timesten_lock_timeout_total	22
Metrics timesten_log_buffer_insert_byte to timesten_remote_update_executes_total	22
Metrics timesten_repl_log_out_busy_total to timesten_repl_transmitter_catch_lag_total	24
Metrics timesten_repl_transmitter_drift_request_total to timesten_table_row_insert_total	25
Metrics timesten_table_row_read_total to timesten_txn_voted_no_total	26

2 Tools for Troubleshooting TimesTen

Using the ttCapture Utility	1
Using the ttIsql Utility	1
Using the ttStatus Utility	2

Using the Logs Generated by the TimesTen Daemon	4
Using the ttTraceMon Utility	4
Starting a Trace and Reading a Trace Buffer	6
SQL Tracing	7
API Tracing	8
DEADLOCK Tracing	9
LOCK Tracing	10
ERR Tracing	12
AGING Tracing	13
AUTOREFRESH Tracing	14
Using the ttXactAdmin Utility	17
About Using SNMP Traps for Diagnostics	18
Using ODBC Tracing	18
Monitoring the TimesTen System Tables	18
Using the Query Optimizer	19

3 Troubleshooting TimesTen Applications and Databases

Unable to Start or Stop TimesTen Daemon	1
No Response from TimesTen Daemon or Subdaemon	2
Check the TimesTen User Error Log	2
Extract a Stack Trace from a Core File	2
Unable to Create Shared Segment	3
Application Unable to Connect to Database in Direct Mode	3
Upgrading Your Database	4
Privileges to Connect to Database	4
Check File System Permissions to Access Database	4
Check the TimesTen Daemon Is Running	4
Check DSN Definition	4
Check DSN Attributes	4
Check Path Name to Database and Transaction Log Directories	4
Manage Semaphores and Shared Memory Segments	5
Check Available Swap Space (Virtual Memory)	5
Increase the Number of Available File Descriptors	6
Troubleshooting Client/Server Problems	6
Cannot Connect to the TimesTen Server	6
TimesTen Server Failed	7
Cannot Find Server DSN	7
TimesTen Server Failed to Load Driver	7
Application Times Out When Accessing TimesTen Server	8
TimesTen Client Loses Connection with TimesTen Server	8
Failed to Attach to Shared Memory Segment for IPC	8

Thread Stack Overflow When Using Multiple Client Connections	8
Out of Space When DSN Specifies New Database	8
Application Connects or Disconnects Are Slow	9
Check If Database Is Being Recovered	9
Check ODBC Tracing	9
Application Is Disconnected Unexpectedly	9
Check for ODBC, JDBC, OCI, Pro*C, and PL/SQL Application Errors	10
Check the User Error Log	10
Application Is Slow	10
Consider Connection Mode	11
Update Statistics for Your Tables	12
Verify Lock and Isolation Levels	12
Check Trace Settings	13
Check Partition Counts for Tables	13
Application Unresponsive, Appears Hung	13
Check Logs and Gather Trace Information	14
Check for ODBC Errors	14
Check for Deadlocks and Timeouts	14
Application Unable to Find Previously Created Objects	14
Check Temporary DSN Attribute	15
Check Overwrite DSN Attribute	15
Check Path Name to Database	15
Troubleshooting OCI and Pro*C/C++ Applications	15
Running Out of a Resource	15
Operating System Tools and Shared Memory	16
Check the Amount of Memory Allocated to a Database	16
Permanent Memory Region Filling Up	17
Temporary Memory Region Filling Up	17
Update Query Optimizer Statistics	17
Check Memory Used by Queries	17
Out of Memory After Fatal Crash of a Database	18
Check Transaction Log File Use of File System Space	18
Check If Tracing Is Enabled	19
Check the Semaphore Limit	19
Duplicate Results from a SELECT Statement	19
Cannot Attach PL/SQL Shared Memory	20

4 Monitoring and Troubleshooting TimesTen Cache

Monitoring Cache Groups with Autorefresh	2
Retrieve Information About Autorefresh Operations	2
Display Information from Change Log Tables	3

Understand Messages About Autorefresh in Support Log	3
Diagnose Autorefresh Failures	4
Diagnose Autorefresh Performance Problems in TimesTen Classic	5
Unable to Create a Cache Group	5
Unable to Start or Stop Cache Agent	6
Check Status of Cache Agent	6
Check ORACLE_HOME Environment Variable	7
Check NLS Environment Variables	7
Unable to Resolve Oracle Service Name	7
Unable to Resolve Connect Identifier	8
Incompatible Oracle Database Server and Client Versions	8
Unable to Validate Oracle Database User Name and Password	8
Check Library Path Environment Variable	9
Check Status of TNS Listener and Oracle Database Server	9
Check Oracle Database Privileges	9
Check Database Source Name Definition	9
Reboot TimesTen Machine	9
Set Cache Administration User Name and Password	10
Check User and System Environment in TimesTen Classic	10
Verify Loaded Dynamic Libraries for TimesTen Classic	10
OCI Initialization Failed	11
Unsupported Data Type Mapping	11
Null Constraint Does Not Match Oracle Database	11
DDL on Cached Oracle Database Tables May Cause Failure	12
Changes Not Visible After Updating Object in Cache Group	12
Load or Refresh Fails	13
Optimize Performance for TimesTen Cache	13
Check Large Batch Jobs on Oracle Database Tables	14
Autorefresh Not Refreshing Cache at Specified Interval	15
Reset Autorefresh State	16
Recover and Reset Autorefresh for Oracle Database Objects	16
Avoid a Lock Timeout Condition When Modifying Autorefresh	17
Incremental Autorefresh Not Progressing	18
Validate Autorefresh for Oracle Database Objects	18
Incremental Autorefresh Becomes Full Autorefresh	19
Detect When Incremental Autorefresh Becomes Full	19
Understanding the Cache Administration User Tablespace	19
Diagnose a Full Cache Administration User Tablespace	20
Monitor Usage of a Cache Administration User Tablespace	20
Considerations When a Tablespace Is Full	21
Poor Autorefresh Performance	22
Unresponsive TimesTen Database Degrades Autorefresh Performance	22

Set Cache TimesTen Database Timeout	23
Configure Recovery Method for Cache Groups	23
Refresh of Autorefresh Cache Group Shows Excessive Wait on Resources	26
Abnormally Large Change Log and Base Tables	27
Fragmented Autorefresh Change Log Table Space	27
Degradation in Performance When Autorefresh Interval Is Small	28
Declaring NOVALIDATE on Constraints Causes Cache Group Failure	28
AWR Report Shows Lock Contention with DBMS_LOCK	30

5 Troubleshooting Replication

Unable to Create a Replication Scheme	1
Unable to Alter a Replication Scheme	2
Unable to Start or Stop Replication Agent	2
Replication Does Not Work	3
Check Status of TimesTen Daemon and Replication Agents	3
Check Replication Agents Are Communicating	5
Check Replication State	6
Check Replication Scheme Configuration	6
Check Owner Names	8
Check Replication Owner	8
Check Table Owner	9
Check Consistency between Replicated Tables	9
Replication Unresponsive, Appears Hung	10
Check Replication State	10
Check Return Receipt Timeout Setting	10
Poor Replication or XLA Performance	10
Check Durability Settings	11
Check Network Bandwidth	11
Check Use of Return Receipt Blocking	11
Check Replication Configuration	11
Check Size of Log Buffer	12
Check Reads from Transaction Log Files	12
Problems Using ttRepAdmin	14
Problems Changing State of Replication Receiver	14
Problems Using ttRepAdmin -duplicate	14
Database Created Before Duplication	14
Wrong Database or Host Name	15
Must Specify Scheme	15
Problems with Conflict Checking	16
Column Invalid for Replication Timestamp	16
Timestamp Does Not Exist	16

Conflict Reporting Slows Down Replication	16
Problems with Parallel Replication	17

6 Troubleshooting AWT Cache Groups

Replication Agent Not Started	1
AWT Performance Monitoring	1
Enable AWT Performance Monitoring	2
Display AWT Performance Results with the ttRepAdmin Utility	2
Use System Tables to Monitor AWT Operations	3
Possible Causes of Poor AWT Performance	3
Problems on the Oracle Database	4
UNLOAD CACHE GROUP Requires Row Updates to Complete Propagation	4
Oracle Database Errors Reported by TimesTen for AWT	4
Permanent Oracle Database Errors Reported by TimesTen	4
Transient Oracle Database Errors Reported by TimesTen	5

About This Content

This document covers TimesTen support for monitoring and troubleshooting operations.

Audience

This document is intended for anyone who wants to monitor and troubleshoot TimesTen databases. It also includes information about the TimesTen Prometheus Exporter.

Documentation Accessibility

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

Access to Oracle Support

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

Conventions

The following text conventions are used in this document.

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

What's New

This section summarizes the new features and functionality of TimesTen Release 22.1 that are documented in this guide, providing links into the guide for more information.

New Features in Release 22.1.1.1.0

- The TimesTen Prometheus Exporter together with Prometheus enables you to monitor the health and operation of your databases. The TimesTen Prometheus Exporter converts TimesTen metrics into the form used by Prometheus. This integration enables you to add TimesTen to the systems that you monitor with Prometheus. See [The TimesTen Prometheus Exporter](#) for details.
- TimesTen supports SNMP version 1, 2, and 3 traps through net-SNMP. See [About Using SNMP Traps for Diagnostics](#) for more information.

1

The TimesTen Prometheus Exporter

The TimesTen Prometheus Exporter (TimesTen exporter or Exporter) together with Prometheus enable you to monitor the health and operation of your databases. The TimesTen exporter converts TimesTen metrics into a form supported by Prometheus. This simple integration enables you to add TimesTen to the systems that you can monitor with Prometheus.

This chapter gives an overview of Prometheus and introduces the Exporter. It explains the function of the `ttExporter` utility. The chapter demonstrates how to configure the Exporter and Prometheus either with client certificate authentication or with no authentication . It also shows you how to configure Prometheus to poll (*scrape*) for metrics. It then shows you how to start and stop the Exporter. The chapter concludes by describing the metrics that the Exporter supports.

Topics:

- [About Prometheus](#)
- [About the TimesTen Exporter](#)
- [About Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication](#)
- [Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication](#)
- [About Configuring the TimesTen Exporter and Prometheus with No Authentication](#)
- [Configuring the TimesTen Exporter and Prometheus with No Authentication](#)
- [The Metrics Supported by the TimesTen Exporter](#)

About Prometheus

Prometheus is an open source systems monitoring and alerting toolkit. It collects and stores metrics from a variety of sources. It has its own time-series database and time-series query language. You can use Prometheus directly to construct near real-time graphs of metrics or to create programmable alerts. For more information on Prometheus, see: <https://prometheus.io/docs/introduction/overview/>

Prometheus consists of a server that runs on your network. For information on downloading and installing Prometheus, see: https://prometheus.io/docs/introduction/first_steps/

You configure your Prometheus server to *scrape* (poll) some set of endpoints for metrics. Typically, Prometheus scrapes for metrics every 15 seconds by sending an HTTP or an HTTPS GET request to each endpoint for the `/metrics` document.

An endpoint provides an HTTP or HTTPS server that responds to a GET on `/metrics` with a plain text ASCII document containing the metrics. Endpoints identify each metric with a name, a descriptive label, and a double-precision floating point value.

For each metric, a Prometheus exporter provides help text and a data type. An Exporter provides this information with lines of text that look like comments. For example, for each metric:

- The first line of text begins with `#HELP` followed by the metric name and the metric description.

- The second line of text begins with `#TYPE` followed by the metric name and the Prometheus metric type. (The metric type is discussed later.)
- The third line of text begins with the metric name, followed by the optional label(s), and then the value of the metric.

Prometheus supports four metric types:

- *counter*: A counter is a cumulative metric that represents a single monotonically increasing counter whose value can only increase or be reset to zero on restart. An example is the number of `SELECT` statements executed. Some counter values reset, such as when you unload and load a TimesTen database from and to memory. Prometheus handles such situations automatically.
- *gauge*: A gauge is a metric that represents a single numerical value that can arbitrarily go up and down. An example is the number of bytes allocated for permanent space for the TimesTen database.
- *histogram*: A histogram samples observations and counts them in configurable buckets. It also provides a sum of all observed values.
- *summary*: A summary samples observations and provides a total count of observations and a sum of all observed values. It calculates configurable quantiles over a sliding time window.

For more information on Prometheus metric types, see: https://prometheus.io/docs/concepts/metric_types/

About the TimesTen Exporter

The TimesTen exporter runs on Linux x86_64 systems. The Exporter is supported in TimesTen Classic and in TimesTen Scaleout.

In TimesTen Scaleout, the Exporter is supported on each host that is running either a data instance or a management instance. If there are multiple data instances on a single host, you deploy one Exporter for each data instance.

The Exporter does not run by default. You must specifically start the Exporter. The Exporter runs independently of the TimesTen instance.

The Exporter consists of a single executable program (utility) called `ttExporter`. When you create a TimesTen instance, the `ttExporter` utility is included as a part of the TimesTen instance. The `ttExporter` utility parses command line arguments and presents itself as an HTTP or an HTTPS server.

When the Exporter receives an HTTP (or HTTPS) GET request to `/metrics`, it retrieves the TimesTen metrics from each TimesTen database that it monitors and prepares a plain text HTTP(S) response with the metrics, including help text and type information.

The TimesTen metrics are collected from a variety of sources, including the TimesTen system tables and views and the TimesTen built-in procedures and utilities. Each TimesTen metric is prefixed with `timesten`. Each TimesTen metric consists of:

- Metric name: For example, `timesten_perm_allocated_bytes`
- Prometheus metric type: Either `counter` or `gauge`
- Description: For example, Bytes of permanent space allocated

See [The Metrics Supported by the TimesTen Exporter](#) for a complete list of all the TimesTen metrics that the Exporter supports.

The Exporter converts each TimesTen metric into a format that Prometheus supports. For example, the Exporter converts the `timesten_perm_allocated_bytes` TimesTen metric into this Prometheus form:

```
# HELP timesten_perm_allocated_bytes Bytes of permanent space allocated
# TYPE timesten_perm_allocated_bytes gauge
timesten_perm_allocated_bytes{dsn="mydatabase",instance="myinstance"} 1073741824
```

A Prometheus exporter can be configured to require client certificate authentication with Transport Layer Security (mutual TLS) or to require no authentication. The TimesTen exporter supports both configurations. While it is typical for a Prometheus exporter to operate with no authentication, the default configuration for the TimesTen exporter is client certificate authentication. In a client certificate authentication configuration, an Oracle wallet is used to store the TLS credentials.

The configuration requirements differ depending on whether you configure the TimesTen exporter with client certificate authentication or with no authentication.

- To configure the Exporter with client certificate authentication, see:
 - [About Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication](#)
 - [Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication](#)
- To configure the Exporter with no authentication, see:
 - [About Configuring the TimesTen Exporter and Prometheus with No Authentication](#)
 - [Configuring the TimesTen Exporter and Prometheus with No Authentication](#)

About Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication

To use the TimesTen exporter and Prometheus with client certificate authentication, you must configure the Exporter and Prometheus to require client certificate authentication.

To configure the Exporter, you must:

- Create a self-signed certificate. The Exporter uses the server certificate to authenticate itself to clients and to authenticate client certificates. You use the `ttExporter` utility to create the self-signed certificate. The self-signed certificate must be created on the same host that the Exporter runs on. The self-signed certificate is stored in an Oracle Wallet in the directory of your choice. If you do not specify a directory, the Oracle Wallet is stored in the user's home (`$HOME`) directory. TimesTen recommends that you do not store the certificate information in the user's `$HOME` directory.
- Export the server certificate. You use the `ttExporter` utility to export the server certificate in PEM format to a file. This file is used by Prometheus for client certificate authentication.
- Create and export both a client certificate and a client private key. You use the `ttExporter` utility to create and export a client certificate to one file and to create and export a client private key to a second file. You must create and export a client certificate and a client private key for each Prometheus instance that scrapes metrics from the Exporter. The client certificate file and the client private key file are used by Prometheus for client certificate authentication.

See ttExporter in the *Oracle TimesTen In-Memory Database Reference* for details on the ttExporter utility.

After you create and export the server certificate, the client certificate, and the client private key, you must configure Prometheus to require client certificate authentication for the Exporter.

Prometheus provides a sample configuration in a file called `prometheus.yml`. This file is available when you download Prometheus. There are two sections (blocks) that you need to modify. The `tls_config` block is used for configuring client certificate authentication. The `scrape_configs` block is used to set the remaining parameters necessary for Prometheus to securely scrape metrics from the Exporter.

Specifically, you must modify these settings in the `tls_config` block:

- `ca_file`: This is the server certificate file that is required by Prometheus for client certificate authentication. Recall that after using the TimesTen ttExporter utility to create the server certificate, you then use the ttExporter utility again to export this server certificate to a file. It is this file that you supply for the Prometheus `ca_file` parameter.
- `cert_file`: This is the client certificate file that is required by Prometheus for client certificate authentication. Recall that you use the TimesTen ttExporter utility to create and export the client certificate to a file. It is this file that you supply for the Prometheus `cert_file` parameter.
- `key_file`: This is the client private key file that is required by Prometheus for client certificate authentication. Recall that you use the TimesTen ttExporter utility to create and export the client private key to a file. It is this file that you supply for the Prometheus `key_file` parameter..

In addition, you must modify these general settings in the `scrape_configs` block:

- `job_name`: This is the job name that you assign to scrape the metrics from the Exporter. You can specify any name for `job_name`.
- `metrics_path`: This is the HTTPS resource path from which Prometheus fetches the TimesTen metrics from the Exporter. The Exporter uses the `/metrics` HTTPS resource path. This is also the Prometheus default HTTPS resource path. Therefore, you do not have to set the `metrics_path` parameter. If not set, Prometheus uses the default `/metrics` for the HTTPS resource path.
- `scheme`: This is the protocol scheme used for requests. The HTTPS protocol is used when you configure the Exporter and Prometheus to require client certificate authentication. You must set the `scheme` parameter to `https`.
- `static_configs` block: This block contains the `targets` parameter. For the `targets` parameter, you specify the host that the Exporter runs on and the port number that the Exporter listens on. The naming convention is `host:portnumber`. For example, if the Exporter runs on the `myhost` host and listens on the `12345` port number, then you set the `targets` parameter to `myhost:12345`.

After you modify the Prometheus configuration file, the next step is to start the Exporter. You use the TimesTen ttExporter utility with the `-port p` option to start the Exporter (where `p` is the port number that the Exporter listens on). This port number must match the port number you specify in the Prometheus configuration file. (Specifically, in the Prometheus configuration file, the port number is the port number you supply to the `targets` parameter of the `static_configs` block.)

Once you start the Exporter, the Exporter continues to run until you stop it with a signal such as `SIGINTR` or `SIGTERM`. You can also have ttExporter write its process ID into a file when you start the Exporter. You use the ttExporter utility with the `-pid-file pidfilename` option to

accomplish this. You can then stop the Exporter by terminating the `ttExporter` process ID that was stored in the `pidfilename` file.

See [Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication](#) for a complete example illustrating how to configure the Exporter and Prometheus to have Prometheus securely scrape TimesTen metrics from the Exporter.

See `ttExporter` in the *Oracle TimesTen In-Memory Database Reference* for details on the `ttExporter` utility.

See <https://prometheus.io/docs/prometheus/latest/configuration/configuration/> and https://prometheus.io/docs/prometheus/latest/configuration/configuration/#scrape_config for more information on the Prometheus configuration file.

Configuring the TimesTen Exporter and Prometheus with Client Certificate Authentication

This example illustrates how to configure both the Exporter and Prometheus to require client certificate authentication. It first uses the `ttExporter` utility to create and export the server certificate, to create and export the client certificate, and to create and export the client private key. It then configures Prometheus to use the server certificate, the client certificate, and the client private key that was created by the TimesTen `ttExporter` utility. The example starts the Exporter. After the Exporter is started, the example verifies that Prometheus is securely scraping the TimesTen metrics that the Exporter exposes.

1. Create a subdirectory to store the Oracle Wallet containing the certificate information used by the Exporter. This example creates the `mycertdir` subdirectory.

```
mkdir -p mycertdir
```

2. Use the `ttExporter` utility with the `-create-server-certificate` to create the self-signed certificate and store an Oracle Wallet containing the certificate information. Use the `-certificate-directory` option to specify the location for the certificate information (`mycertdir`, in this example). If you do not specify `-certificate-directory`, `ttExporter` creates and stores an Oracle Wallet containing the certificate information in your user's `$HOME` directory. Perform this step one time.

```
% ttExporter -create-server-certificate -certificate-directory mycertdir
% ls -a mycertdir
.  ..  .ttwallet.C76FCD9A-E5C8-4AC5-A73D-7DAB72203E43
% ls -a mycertdir/.ttwallet.C76FCD9A-E5C8-4AC5-A73D-7DAB72203E43
.  ..  cwallet.sso
```

3. Use the `ttExporter` utility with the `-export-server-certificate` option to export the server certificate in PEM format. This example exports the server certificate to the `server.crt` file in the `mycertdir` directory. You must provide the `-certificate-directory` option since you specified it when you created the server certificate. This example also verifies that the `mycertdir/server.crt` file was created. Note that you will later supply the `mycertdir/server.crt` file for the `ca_file` parameter in the `tls_config` block of the Prometheus configuration file.

```
% ttExporter -export-server-certificate mycertdir/server.crt -certificate-
directory mycertdir
% cat mycertdir/server.crt
-----BEGIN CERTIFICATE-----
```

```

MIIDtzCCAp+gAwIBAgIJANV4RZyydnFrMA0GCSqGSIb3DQEBCwUAMHIxCzAJBgNV
...
4PtJlBCXt6ogsceJI8xrCxcfUjPh9TxBWQfeW2vI+AfmUEzXTrYa+xSmkg==
-----END CERTIFICATE-----

```

4. Use the `ttExporter` utility with the `-export-client-certificate` and the `-export-client-private-key` options to create and export a client certificate to a file (`mycertdir/client.crt`, in this example) and a client private key to a second file (`mycertdir/key.crt`, in this example). You must also provide the `-certificate-directory` option, since you specified it when you created the server certificate. These options must be specified together. This step also verifies that the `mycertdir/client.crt` and the `mycertdir/key.crt` were created. Note that you will later supply the `mycertdir/client.crt` file for the `cert_file` parameter and the `mycertdir/key.crt` file for the `key_file` in the `tls_config` block of the Prometheus configuration file.

 **Note**

Create a client certificate and private key for each Prometheus instance that scrapes metrics from the Exporter.

```

% ttExporter -export-client-certificate mycertdir/client.crt
      -export-client-private-key mycertdir/key.crt -certificate-directory
mycertdir
% cat mycertdir/client.crt
-----BEGIN CERTIFICATE-----
MIIDXTCCAkUCCCPY0o/eJJFMMMA0GCSqGSIb3DQEBCwUAMHIxCzAJBgNVBAYTA1VT
...
c18fFPkKcsJIJqYudoY8ulmIzThdjVhA8zY25vPU6exxAKw1BmJPuwqexyIpbNNH
LA==
-----END CERTIFICATE-----
% cat mycertdir/key.crt
-----BEGIN RSA PRIVATE KEY-----
MIIEpAIBAAKCAQEA0/+ERz/MgA34Uv+UNn1RPDnss14QKyevd05vc/1wrSDwEZJy
...
1+Mi8MgmjKInXkmRotxWbJ6LSQY5wiol5HhwLTTCQgoCdVTO+usaXQ==
-----END RSA PRIVATE KEY-----

```

5. Configure Prometheus to securely scrape metrics from the Exporter. This requires that you modify the Prometheus configuration file (`prometheus.yml`) by setting the `ca_file`, the `cert_file`, and the `key_file` parameters in the `tls_config` block of the `scrape` configuration (`scrape_configs`) block. You also must set general parameter settings in the `scrape_configs` block.

Specifically, modify these parameters in the `scrape_configs` block:

- `job_name`: Specify the job name assigned to scrape the metrics (`tt_12345`, in this example).
- `metrics_path`: There is no need to specify the `metrics_path` as the default is `/metrics`. The Exporter uses the `/metrics` default as the HTTPS resource path on which Prometheus fetches the TimesTen metrics from the Exporter.
- `scheme`: You must specify the `https` scheme to ensure the Exporter and the Prometheus instance operate with client server authentication.
- `tls_config` block:

- ca_file: Specify the server certificate file that was exported when you ran `ttExporter -export-server-certificate` (`mycertdir/server.crt`, in this example).
- cert_file: Specify the client certificate file that was created and exported when you ran `ttExporter -export-client-certificate` (`mycertdir/client.crt`, in this example)
- key_file: Specify the client private key that was created and exported when you ran `ttExporter - export-client-private-key` (`mycertdir/key.crt`, in this example)
- static_configs block: For the targets parameter, specify the host that the Exporter will be running on and the port number that the Exporter will be listening on (`myhost:12345`, in this example).

```
vi prometheus.yml
...
scrape_configs:
  - job_name: 'tt_12345'
    # metrics_path defaults to '/metrics'
    # scheme defaults to 'http'.
    scheme: https
    tls_config:
      ca_file: mycertdir/server.crt
      cert_file: mycertdir/client.crt
      key_file: mycertdir/key.crt
    static_configs:
      - targets: ['myhost:12345']
```

6. Use the `ttExporter` utility to start the Exporter. You must specify the `-port` option. Ensure to specify the same port as what was specified in the Prometheus configuration file (12345, in this example). You must also provide the `-certificate-directory` option since this option was specified when you create the server certificate. Optionally, specify the `-pid-file` option to store the process ID of `ttExporter` in a file. In this example, the file is `/tmp/ttexporter.pid` and the process ID is 24302.

```
% ttExporter -port 12345 -pid-file /tmp/ttexporter.pid -certificate-
directory mycertdir
% cat /tmp/ttexporter.pid
24302
% ps 24302
  PID TTY      STAT      TIME COMMAND
24302 pts/0    Sl+      0:00 timesten_home/bin/_ttExporter -port 12345 -pid-
file /tmp/ttexporter.pid
-certificate-directory mycertdir
```

7. Optional and for verification purposes only: Navigate to the certificate directory (`mycertdir`, in this example). Then use the Linux `curl` command to verify the TimesTen metrics are exposed and available for Prometheus to scrape.

```
% cd mycertdir
% curl --cacert server.crt --cert ./client.crt --key ./key.crt https://
myhost:12345/metrics
...
# HELP timesten_perm_in_use_bytes Bytes of permanent space used
# TYPE timesten_perm_in_use_bytes gauge
```

```
timesten_perm_in_use_bytes{dsn="mydsn",instancename="myinstance"} 20385792
# HELP timesten_perm_in_use_high_water_bytes Maximum bytes of permanent
space used
# TYPE timesten_perm_in_use_high_water_bytes gauge
timesten_perm_in_use_high_water_bytes{dsn="mydsn",instancename="myinstance"
} 20385792
...
```

8. Optional: This step illustrates that you can stop the Exporter. This terminates the ttExporter process. Prometheus no longer scrapes the endpoint. Recall that the -pid-file option was specified when you started the Exporter. Use the Linux pkill -F command to terminate the process ID. Recall the example stores the process ID in the /tmp/ttexporter.pid file and that the process ID is 24302.

```
% pkill -F /tmp/ttexporter.pid
% cat /tmp/ttexporter.pid
24302
% ps 24302
  PID TTY      STAT   TIME COMMAND
```

You successfully configured the Exporter and Prometheus to require client certificate authentication. You also started the Exporter. Prometheus is securely scraping the TimesTen metrics that the Exporter exposes at the /metrics endpoint. Prometheus continues to scrape the TimesTen metrics until you stop the Exporter.

About Configuring the TimesTen Exporter and Prometheus with No Authentication

To use the TimesTen exporter and Prometheus without authentication, you must configure Prometheus for use without authentication and then start the Exporter (without authentication).

Prometheus provides a sample configuration in a file called `prometheus.yml`. This file is available when you download Prometheus. In this file is a section (block) called `scrape_configs`. You modify the settings of the parameters in this block in order to configure Prometheus to scrape metrics from the Exporter and to scrape such metrics without authentication.

In particular, you modify these settings in the `scrape_configs` block:

- `job_name`: This is the job name that you assign to scrape the metrics from the Exporter. You can specify any name for `job_name`.
- `metrics_path`: This is the HTTP resource path from which Prometheus fetches the TimesTen metrics from the Exporter. The Exporter uses the `/metrics` HTTP resource path. This is also the Prometheus default HTTP resource path. Therefore, you do not have to set the `metrics_path` parameter. If not set, Prometheus uses the default `/metrics` for the HTTP resource path.
- `scheme`: This is the protocol scheme used for requests. The HTTP protocol is used when you configure the Exporter and Prometheus without authentication. This is also the Prometheus default. Therefore, you do not have to set the `scheme` parameter.
- `static_configs` block: This block contains the `targets` parameter. For the `targets` parameter, you specify the host that the Exporter runs on and the port number that the Exporter listens on. The naming convention is `host:portnumber`. For example, if the

Exporter runs on the `myhost` host and listens on the 12346 port number, then you set the `targets` parameter to `myhost:12346`.

After you modify the Prometheus configuration file, the next step is to start the Exporter. You use the TimesTen `ttExporter` utility with the `-insecure` option to start the Exporter. The `-insecure` option denotes no authentication. The `-port p` option is required when you start the Exporter (where `p` is the port number that the Exporter listens on). This port number must match the port number you specify in the Prometheus configuration file. (Specifically, in the Prometheus configuration file, the port number is the port number you supply to the `targets` parameter of the `static_configs` block.)

Once you start the Exporter, the Exporter continues to run until you stop it with a signal such as `SIGINTR` or `SIGTERM`. You can also have `ttExporter` write its process ID into a file when you start the Exporter. You use the `ttExporter` utility with the `-pid-file pidfilename` option to accomplish this. You can then stop the Exporter by terminating the `ttExporter` process ID that was stored in the `pidfilename` file.

See [Configuring the TimesTen Exporter and Prometheus with No Authentication](#) for a complete example illustrating how to configure the Exporter and Prometheus to have Prometheus scrape TimesTen metrics from the Exporter (without authentication).

See `ttExporter` in the *Oracle TimesTen In-Memory Database Reference* for details on the `ttExporter` utility.

See <https://prometheus.io/docs/prometheus/latest/configuration/configuration/> and https://prometheus.io/docs/prometheus/latest/configuration/configuration/#scrape_config for more information on the Prometheus configuration file.

Configuring the TimesTen Exporter and Prometheus with No Authentication

This example first illustrates how to configure Prometheus without authentication. It then uses the `ttExporter` utility to start the TimesTen exporter. The `-insecure` option of the `ttExporter` utility is used to denote no authentication. After Prometheus is configured and the Exporter is started, Prometheus can then scrape the TimesTen metrics that the Exporter exposes.

1. Configure Prometheus to scrape metrics from the TimesTen exporter by modifying the `scrape_config` block of the Prometheus configuration file (`prometheus.yml`, in this example).

Modify these parameters in the `scrape_configs` block of the Prometheus configuration file (`prometheus.yml`, in this example):

- `job_name`: Specify the job name assigned to scrape the metrics (`tt_12346`, in this example).
- `metrics_path`: The Exporter uses `/metrics` as the HTTP resource path, which is also the Prometheus default. Therefore, you do not need to set the `metrics_path` parameter.
- `scheme`: The Exporter uses the HTTP protocol scheme which is the scheme that is used for no authentication. This is also the Prometheus default. Therefore, you do not need to set the `scheme` parameter.
- `static_configs` section: You must specify the name of the host and the port number for the `targets` parameter. The host is the server that Exporter will run on and the port number is the port that the Exporter will listen on (`myhost:12346`, in this example).

ⓘ Note

The port number must be the same as the port number you use to start the Exporter.

```
vi prometheus.yml
...
scrape_configs:
- job_name: 'tt_12346'
  # metrics_path defaults to '/metrics'
  # scheme defaults to 'http'.
  static_configs:
    - targets: ['myhost:12346']
...

```

2. Use the `ttExporter` utility to start the Exporter. Specify the `-insecure` and `-port` options. Both of these options are required. Ensure to specify the same port as what was specified in the Prometheus configuration file (12346, in this example). Optionally, specify the `-pid-file` option to store the process ID of `ttExporter` in a file. In this example, the file is `/tmp/ttexporter.pid2` and the process ID is 942.

```
% ttExporter -insecure -port 12346 -pid-file /tmp/ttExporter2.pid
% cat /tmp/ttexporter2.pid
942
```

3. Optional: From your browser, enter `http://myhost:12346/metrics`. This enables you to verify that the Exporter has converted the TimesTen metrics into the form required by Prometheus. See [The Metrics Supported by the TimesTen Exporter](#) for the supported TimesTen metrics.

```
In browser: http://myhost:12346/metrics
...
# HELP timesten_perm_in_use_bytes Bytes of permanent space used
# TYPE timesten_perm_in_use_bytes gauge
timesten_perm_in_use_bytes{dsn="mydsn",instancename="myinstance"} 20385792
# HELP timesten_perm_in_use_high_water_bytes Maximum bytes of permanent
space used
# TYPE timesten_perm_in_use_high_water_bytes gauge
timesten_perm_in_use_high_water_bytes{dsn="mydsn",instancename="myinstance"} 20385792
...

```

4. Optional: This step illustrates that you can stop the Exporter. This terminates the `ttExporter` process. Prometheus no longer scrapes the endpoint. Recall that the `-pid-file` option was specified when you started the Exporter. Use the Linux `kill -F` command to terminate the process ID. Recall the example stores the process ID in the `/tmp/ttexporter2.pid` file and that the process ID is 942.

```
% pkill -F /tmp/ttexporter2.pid
% cat /tmp/ttexporter2.pid
942
% ps 942
  PID  TTY      STAT      TIME  COMMAND
```

You successfully configured Prometheus and started the Exporter. Prometheus is scraping the TimesTen metrics that the Exporter exposes at the `/metrics` endpoint. Prometheus continues to scrape the TimesTen metrics until you stop the Exporter.

The Metrics Supported by the TimesTen Exporter

The TimesTen exporter supports the following metrics:

- TimesTen version
- Contents of the `SYS.SYSTEMSTATS` system table
- Contents of the `SYS.MONITOR` system table
- Contents of the `SYS.V$HEAP_INFO` system view
- Contents of the `SYS.V$GRIDSTATS` system view (for TimesTen Scaleout)
- Contents of the `SYS.V$BOOKMARK` system view
- Contents of the `SYS.V$CKPT_HISTORY` system view
- Results from the `ttLogHolds` built-in procedure
- Results from the `ttPLSQLMemoryStats` built-in procedure
- Results from the `ttLatchStatsGet` built-in procedure
- Results from the `ttStatus` and the `ttGridAdmin` utilities
- Operating system thread statistics for the TimesTen main daemon and sub-daemon threads
- Total size and allocated size (in bytes) of the file system space for the checkpoint files and the transaction log directories
- Metrics about the TimesTen exporter itself, including
 - Current memory consumption
 - Time to service the current metric request
 - Number of metric requests
 - Number of disallowed requests
 - Total time to service metric requests
 - Version of the exporter

See Overview of System Tables and Views in the *Oracle TimesTen In-Memory Database System Tables and Views Reference* for information on the TimesTen system tables and system views. See Built-In Procedures and Utilities for information on the TimesTen built-in procedures and the TimesTen utilities.

Some of the metrics may not return any information, depending on the configuration. Some of the metrics pertain only to TimesTen Classic and some only to TimesTen Scaleout. Note that with few exceptions, the TimesTen exporter does not publish metrics for the underlying operating system. You can use the *Node Exporter* distributed with Prometheus to collect that information. For information on the Prometheus Node Exporter, see <https://prometheus.io/docs/guides/node-exporter/>.

The following tables show the metric name, the metric type, and the metric description for each TimesTen exporter metric. The metrics are arranged in alphabetical order and expand across multiple tables:

- [Metrics `timesten_aging_lru_commit_total` to `timesten_cg_autorefresh_row_update_total`](#)
- [Metrics `timesten_cg_awt_oracle_call_total` to `timesten_cg_flushed_row_total`](#)
- [Metrics `timesten_cg_parawt_auto_apply_applctn_seqnum` to `timesten_cg_parawt_auto_apply_partitions`](#)
- [Metrics `timesten_cg_parawt_auto_apply_plsqlmode_batches` to `timesten_cg_parawt_auto_apply_validbatch_waits`](#)
- [Metrics `timesten_cg_parawt_auto_rx_batches` to `timesten_cg_parawt_auto_sorter_txns_waits`](#)
- [Metrics `timesten_cg_parawt_auto_tx_applier_waits` to `timesten_ckpt_write_total`](#)
- [Metrics `timesten_collections_timedout` to `timesten_cs_server_transmitted_byte_total`](#)
- [Metrics `timesten_cursor_close_total` to `timesten_index_range_rebuild_insert_total`](#)
- [Metrics `timesten_index_range_row_fetch_repl_total` to `timesten_lock_timeout_total`](#)
- [Metrics `timesten_log_buffer_insert_byte` to `timesten_remote_update_executes_total`](#)
- [Metrics `timesten_repl_log_out_busy_total` to `timesten_repl_transmitter_catch_lag_total`](#)
- [Metrics `timesten_repl_transmitter_drift_request_total` to `timesten_table_row_insert_total`](#)
- [Metrics `timesten_table_row_read_total` to `timesten_txn_voted_no_total`](#)

Metrics `timesten_aging_lru_commit_total` to `timesten_cg_autorefresh_row_update_total`

Metric	Type	Description
<code>timesten_aging_lru_commit_total</code>	counter	Number of LRU aging commits done since the database was loaded into memory
<code>timesten_aging_lru_cycle_total</code>	counter	Number of LRU aging cycles completed since the database was loaded into memory
<code>timesten_aging_lru_high_threshold_reached_total</code>	counter	Number of times LRU aging high threshold is reached since the database was loaded into memory
<code>timesten_aging_lru_low_threshold_reached_total</code>	counter	Number of times LRU aging low threshold is reached since the database was loaded into memory
<code>timesten_aging_lru_row_deleted_total</code>	counter	Number of rows deleted during LRU aging since the database was loaded into memory
<code>timesten_aging_lru_row_skipped_total</code>	counter	Number of rows that were not deleted using LRU aging because of lock contention since the database was loaded into memory
<code>timesten_aging_timebased_commit_total</code>	counter	Number of time-based aging commits done since the database was loaded into memory
<code>timesten_aging_timebased_cycle_total</code>	counter	Number of time-based aging cycles completed since the database was loaded into memory
<code>timesten_aging_timebased_row_deleted_total</code>	counter	Number of rows deleted during time-based aging since the database was loaded into memory

Metric	Type	Description
timestamp_aging_timebased_row_skipped_total	counter	Number of rows that were not deleted using time-based aging because of lock contention since the database was loaded into memory
timestamp_cg_ar_flush_sync_total	counter	Number of times flush and sync had been called because cache config CacheCommitDurable is set to 0
timestamp_cg_ar_nondurable_commit_total	counter	Number of times non-durable commit was used during refresh
timestamp_cg_autorefresh_cycle_total	counter	Number of autorefresh cycles completed successfully on TimesTen
timestamp_cg_autorefresh_failed_cycle_total	counter	Number of autorefresh cycles that failed because of errors
timestamp_cg_autorefresh_full_refresh_total	counter	Number of full refreshes triggered during autorefresh operations
timestamp_cg_autorefresh_logtblspacepct_exceed_ed_total	counter	Number of iterations the cache agent has attempted to garbage collect
timestamp_cg_autorefresh_row_delete_total	counter	Number of rows deleted in TimesTen during autorefresh from the Oracle database
timestamp_cg_autorefresh_row_insert_total	counter	Number of rows inserted in TimesTen during autorefresh from the Oracle database
timestamp_cg_autorefresh_row_update_total	counter	Number of rows updated in TimesTen during autorefresh from the Oracle database

Metrics timestamp_cg_awt_oracle_call_total to timestamp_cg_flushed_row_total

Metric	Type	Description
timestamp_cg_awt_oracle_call_total	counter	Number of calls made to the Oracle database
timestamp_cg_awt_oracle_commit_total	counter	Number of AWT transactions committed on the Oracle database
timestamp_cg_awt_oracle_retry_total	counter	Number of times AWT transactions are retried on the Oracle database in case of an error
timestamp_cg_awt_oracle_rollback_total	counter	Number of rollbacks on the Oracle database because of errors
timestamp_cg_awt_plsql_mode_batch_total	counter	Number of PL/SQL block batches sent to the Oracle database (cacheawtmethod = 1)
timestamp_cg_awt_plsql_mode_byte_total	counter	Number of bytes sent to the Oracle database in PL/SQL mode (cacheawtmethod = 1)
timestamp_cg_awt_plsql_mode_row_delete_total	counter	Number of rows deleted on the Oracle database in PL/SQL mode (cacheawtmethod = 1)
timestamp_cg_awt_plsql_mode_row_insert_total	counter	Number of rows inserted on the Oracle database in PL/SQL mode (cacheawtmethod = 1)

Metric	Type	Description
timestamp_cg_awt_plsql_mode_row_update_total	counter	Number of rows updated on the Oracle database in PL/SQL mode (cacheawtmethod = 1)
timestamp_cg_awt_sql_mode_batch_total	counter	Number of batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_byte_total	counter	Number of bytes sent to the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_delete_batch_total	counter	Number of delete batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_insert_batch_total	counter	Number of insert batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_row_delete_total	counter	Number of rows deleted on the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_row_insert_total	counter	Number of rows inserted on the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_row_update_total	counter	Number of rows updated on the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_sql_mode_update_batch_total	counter	Number of update batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
timestamp_cg_awt_tt_txn_total	counter	Number of TimesTen transactions propagated to the Oracle database
timestamp_cg_dynamic_local_hit_total	counter	Dynamic cache group cache hits: Number of dynamic loads that find the requested data within the database
timestamp_cg_dynamic_local_miss_logoff_retry_total	counter	Number of times the LogOff had to be retried due to no connections being available when UseCacheConnPool and ConnNoWait are enabled
timestamp_cg_dynamic_local_miss_oracle_load_disconnect_total	counter	Number of times the connection to Oracle was closed after performing a dynamic load
timestamp_cg_dynamic_local_miss_oracle_load_misses_total	counter	Number of data misses on the Oracle database when servicing dynamic load misses for dynamic global cache groups
timestamp_cg_dynamic_local_miss_oracle_load_successful_total	counter	Number of data loads from the Oracle database when servicing dynamic load misses for dynamic cache groups
timestamp_cg_dynamic_local_miss_oracle_load_to_attempt_total	counter	Number of data load attempts from the Oracle database when servicing dynamic load misses for dynamic cache groups
timestamp_cg_dynamic_local_miss_total	counter	Dynamic cache group cache misses: Number of dynamic loads that do not find the requested data within the database and need to load the data from the Oracle database

Metric	Type	Description
<code>timesten_cg_flush_total</code>	counter	Number of flush cache group executions
<code>timesten_cg_flushed_byte_total</code>	counter	Number of bytes flushed to the Oracle database
<code>timesten_cg_flushed_row_total</code>	counter	Number of rows flushed to the Oracle database

Metrics `timesten_cg_parawt_auto_apply_applctn_seqnum` to `timesten_cg_parawt_auto_apply_partitions`

Metric	Type	Description
<code>timesten_cg_parawt_auto_apply_applctn_seqnum</code>	gauge	seqnum of applied ctn
<code>timesten_cg_parawt_auto_apply_applctn_timesta</code>	gauge	timestamp of applied ctn
<code>timesten_cg_parawt_auto_apply_batch_first_seq</code>	gauge	Sequence component of CTN belonging to the first transaction in a parallel AWT batch
<code>timesten_cg_parawt_auto_apply_batch_first_tim</code>	gauge	Timestamp component of CTN belonging to the first transaction in a parallel AWT batch
<code>timesten_cg_parawt_auto_apply_batch_id</code>	gauge	Batch id that the applier is executing
<code>timesten_cg_parawt_auto_apply_batch_last_seqn</code>	gauge	Sequence component of CTN belonging to the last transaction in a parallel AWT batch
<code>timesten_cg_parawt_auto_apply_batch_last_time</code>	gauge	Timestamp component of CTN belonging to the last transaction in a parallel AWT batch
<code>timesten_cg_parawt_auto_apply_batches_beginwa</code>	counter	Number of times begin dependencies the applier waited on
<code>timesten_cg_parawt_auto_apply_batches_full</code>	counter	Number of batches propagated to the Oracle database
<code>timesten_cg_parawt_auto_apply_batches_partial</code>	counter	Number of batches propagated to the Oracle database
<code>timesten_cg_parawt_auto_apply_depend_begin_ba</code>	counter	Batch id that the applier is waiting for tchId
<code>timesten_cg_parawt_auto_apply_depend_begin_se</code>	gauge	CTN seqnum that the applier is waiting for qnum
<code>timesten_cg_parawt_auto_apply_depend_begin_ti</code>	gauge	CTN timestamp that the applier is waiting for mestamp
<code>timesten_cg_parawt_auto_apply_depend_begin_wa</code>	counter	Number of times applier waits on begin dependency to be fulfilled
<code>timesten_cg_parawt_auto_apply_depend_begin_wa</code>	counter	Number of times applier waits on begin dependency to be fulfilled
<code>timesten_cg_parawt_auto_apply_depend_commit_b</code>	counter	Batch id that the applier is waiting for attachId
<code>timesten_cg_parawt_auto_apply_depend_commit_w</code>	counter	Number of times applier waits on commit dependency to be fulfilled
<code>timesten_cg_parawt_auto_apply_depend_commit_w</code>	counter	Number of times applier waits on commit dependency to be fulfilled

Metric	Type	Description
<code>timesten_cg_parawt_auto_apply_oracle_calls</code>	counter	Number of calls made to the Oracle database
<code>timesten_cg_parawt_auto_apply_oracle_commits</code>	counter	Number of AWT transactions committed on the Oracle database
<code>timesten_cg_parawt_auto_apply_oracle_retries</code>	counter	Number of times AWT transactions are retried on the Oracle database in case of an error
<code>timesten_cg_parawt_auto_apply_oracle_rollback_s</code>	counter	Number of rollbacks on the Oracle database because of errors
<code>timesten_cg_parawt_auto_apply_oracle_rollback_s_lock</code>	counter	Number of rollbacks on the Oracle database because some one is waiting for the locks held by this batch
<code>timesten_cg_parawt_auto_apply_oracle_rollback_s_timeout</code>	counter	Number of rollbacks on the Oracle database because we timed out waiting for commit
<code>timesten_cg_parawt_auto_apply_partitions</code>	counter	Number of partitions processed by the applier

Metrics `timesten_cg_parawt_auto_apply_plsqlmode_batches` to `timesten_cg_parawt_auto_apply_validbatch_waits`

Metric	Type	Description
<code>timesten_cg_parawt_auto_apply_plsqlmode_batches</code>	counter	Number of PL/SQL block batches sent to the Oracle database (cacheawtmethod = 1)
<code>timesten_cg_parawt_auto_apply_plsqlmode_bytes</code>	counter	Number of bytes sent to the Oracle database in PL/SQL mode (cacheawtmethod = 1)
<code>timesten_cg_parawt_auto_apply_plsqlmode_deletes_rows</code>	counter	Number of rows deleted on the Oracle database in PL/SQL mode (cacheawtmethod = 1)
<code>timesten_cg_parawt_auto_apply_plsqlmode_inserts_rows</code>	counter	Number of rows inserted on the Oracle database in PL/SQL mode (cacheawtmethod = 1)
<code>timesten_cg_parawt_auto_apply_plsqlmode_switches</code>	counter	Number of times cacheawtmethod is internally switched to handle edge cases
<code>timesten_cg_parawt_auto_apply_plsqlmode_time</code>	gauge	Time spent for the Oracle database accesses for AWT in PL/SQL mode
<code>timesten_cg_parawt_auto_apply_plsqlmode_updates_rows</code>	counter	Number of rows updated on the Oracle database in PL/SQL mode (cacheawtmethod = 1)
<code>timesten_cg_parawt_auto_apply_queue_waits</code>	counter	Number of times applier waits on sorter to fill the partition queue
<code>timesten_cg_parawt_auto_apply_sqlmode_batches</code>	counter	Number of batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_sqlmode_bytes</code>	counter	Number of bytes sent to the Oracle database in SQL mode (cacheawtmethod = 0)

Metric	Type	Description
<code>timesten_cg_parawt_auto_apply_sqlmode_deletes_batches</code>	counter	Number of delete batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_sqlmode_deletes_rows</code>	counter	Number of rows deleted on the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_sqlmode_inserts_batches</code>	counter	Number of insert batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_sqlmode_inserts_rows</code>	counter	Number of rows inserted on the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_sqlmode_time</code>	gauge	Time spent for the Oracle databases accesses for AWT in SQL mode
<code>timesten_cg_parawt_auto_apply_sqlmode_updates_batches</code>	counter	Number of update batches sent to the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_sqlmode_updates_rows</code>	counter	Number of rows updated on the Oracle database in SQL mode (cacheawtmethod = 0)
<code>timesten_cg_parawt_auto_apply_state</code>	gauge	Current state of the applier
<code>timesten_cg_parawt_auto_apply_time_commit</code>	counter	Time spent in Oracle database commits
<code>timesten_cg_parawt_auto_apply_time_total</code>	counter	Total time spent in applying AWT buffers
<code>timesten_cg_parawt_auto_apply_txns_bl</code>	counter	Number of TimesTen transactions not relevant for AWT propagation
<code>timesten_cg_parawt_auto_apply_txns_processed</code>	counter	Number of TimesTen transactions propagated to the Oracle database
<code>timesten_cg_parawt_auto_apply_txns_validbatch_waits</code>	counter	Number of transactions with valid begin dependencies
<code>timesten_cg_parawt_auto_apply_validbatch_dist</code>	counter	Total distance for begin dependencies
<code>timesten_cg_parawt_auto_apply_validbatch_wait_s</code>	counter	Number of batches with begin dependencies

Metrics `timesten_cg_parawt_auto_rx_batches` to `timesten_cg_parawt_auto_sorter_txns_waits`

Metric	Type	Description
<code>timesten_cg_parawt_auto_rx_batches</code>	counter	Number of receiver batches
<code>timesten_cg_parawt_auto_rx_head_ctn_seqnum</code>	gauge	Sequence number of current head's CTN by this receiver into sorter's queue
<code>timesten_cg_parawt_auto_rx_head_ctn_timestamp</code>	gauge	Timestamp of current head's CTN by this receiver into sorter's queue
<code>timesten_cg_parawt_auto_rx_head_prevctn_seqnum</code>	gauge	Sequence number of previous CTN of current head by this receiver into sorter's queue

Metric	Type	Description
timesten_cg_parawt_auto_rx_head_prevctn_timesamp	gauge	Timestamp of previous CTN of current head by this receiver into sorter's queue
timesten_cg_parawt_auto_rx_lastenq_ctn_seqnum	gauge	Sequence number of last CTN enqueued by this receiver into sorter's queue
timesten_cg_parawt_auto_rx_lastenq_ctn_timestamp	gauge	Timestamp of last CTN enqueued by this receiver into sorter's queue
timesten_cg_parawt_auto_rx_skips	counter	Number of transactions thrown away because of unexpected CTN
timesten_cg_parawt_auto_rx_state	gauge	Current state of the receiver
timesten_cg_parawt_auto_rx_time_batch	counter	Time spent in preparing AWT buffers
timesten_cg_parawt_auto_rx_time_enq	counter	Time spent in enqueuing parallel awt txns
timesten_cg_parawt_auto_rx_time_txngen	counter	Time spent in generating parallel awt txns
timesten_cg_parawt_auto_rx_txns	counter	Number of received txns (including blanks)
timesten_cg_parawt_auto_sorter_applctn_seqnum	gauge	seqnum of applied ctn sorter
timesten_cg_parawt_auto_sorter_applctn_timestamp	gauge	timestamp of applied ctn sorter
timesten_cg_parawt_auto_sorter_applfctn_seqnum	gauge	seqnum of applied fctn sorter
timesten_cg_parawt_auto_sorter_applfctn_timestamp	gauge	timestamp of applied fctn sorter
timesten_cg_parawt_auto_sorter_partitions_free	counter	Number of free partitions that can be used by the sorter
timesten_cg_parawt_auto_sorter_partitions_processed	counter	Number of partitions dispatched by the sorter
timesten_cg_parawt_auto_sorter_queue_waits	counter	Number of times the sorter has to wait for the queue to give space
timesten_cg_parawt_auto_sorter_state	gauge	Current state of the sorter
timesten_cg_parawt_auto_sorter_tailctn_seqnum	gauge	seqnum of tail ctn sorter
timesten_cg_parawt_auto_sorter_tailctn_timestamp	gauge	timestamp of tail ctn sorter
timesten_cg_parawt_auto_sorter_tailfctn_seqnum	gauge	seqnum of tail fctn sorter
timesten_cg_parawt_auto_sorter_tailfctn_timestamp	gauge	timestamp of tail fctn sorter
timesten_cg_parawt_auto_sorter_txns_blank	counter	Number of TimesTen transactions not relevant for AWT propagation
timesten_cg_parawt_auto_sorter_txns_processed	counter	Number of transactions processed by the sorter
timesten_cg_parawt_auto_sorter_txns_waits	counter	Number of times the sorter has to wait for appropriate transaction

Metrics timesten_cg_parawt_auto_tx_applier_waits to timesten_ckpt_write_total

Metric	Type	Description
timesten_cg_parawt_auto_tx_applier_waits	counter	Number of times transmitter has to wait before it can transmit again
timesten_cg_parawt_auto_tx_applier_waits_timeouts	counter	Number of times transmitter has to wait before it can transmit again
timesten_cg_parawt_auto_tx_batches	counter	Number of transaction batches sent by the transmitter
timesten_cg_parawt_auto_tx_master_oracle_connектs	counter	Number of attempts by master transmitter to connect to Oracle
timesten_cg_parawt_auto_tx_master_restarts	counter	Number of times master transmitter restarts
timesten_cg_parawt_auto_tx_state	gauge	Current state of the transmitter
timesten_cg_parawt_auto_tx_txns	counter	Number of transactions sent by the transmitter
timesten_cg_swt_row_delete_total	counter	Number of rows in SWT cache groups deleted on the Oracle database
timesten_cg_swt_row_insert_total	counter	Number of rows in SWT cache groups inserted on the Oracle database
timesten_cg_swt_row_update_total	counter	Number of rows in SWT cache groups updated on the Oracle database
timesten_channel_invalidations_total	counter	Number of channel invalidations
timesten_channel_recv_bytes_total	counter	Number of bytes received
timesten_channel_recv_messages_total	counter	Number of messages received
timesten_channel_send_bytes_total	counter	Number of bytes sent
timesten_channel_send_messages_total	counter	Number of messages sent
timesten_checkpoint_avail_bytes	gauge	Available bytes on file system holding checkpoints
timesten_checkpoint_block_written_total	counter	Number of blocks written to checkpoint files
timesten_checkpoint_byte_written_total	counter	Number of bytes written to checkpoint files
timesten_checkpoint_fuzzy_total	counter	Number of fuzzy checkpoint operations
timesten_checkpoint_size_bytes	gauge	Total size in bytes of file system holding checkpoints
timesten_checkpoint_total	counter	Number of checkpoint operations
timesten_checkpoint_write_total	counter	Number of writes to checkpoint files
timesten_ckpt_completed_total	counter	Number of checkpoints completed
timesten_ckpt_fuzzy_completed_total	counter	Number of fuzzy checkpoints completed
timesten_ckpt_recovery_write_byte_total	counter	Number of bytes written for checkpointing during database recovery
timesten_ckpt_write_total	counter	Number of checkpoint writes

Metrics timesten_collections_timedout to timesten_cs_server_transmitted_byte_total

Metric	Type	Description
timesten_collections_timedout	gauge	Database collections that timed out
timesten_commit_buf_overflow_total	counter	Number of commits that overflowed the buffer
timesten_commit_dml_on_success_total	counter	Number of DML commit on success that can be bundled with execution
timesten_compact_total	counter	Number of compaction operations
timesten_connect_total	counter	Number of connections established
timesten_connection_cs_established	gauge	Number of client/server connections established
timesten_connection_disconnected	gauge	Number of database disconnects
timesten_connection_established	gauge	Number of database connections established
timesten_connection_established_direct	gauge	Number of direct-linked database connections established
timesten_connection_established_first_total	counter	Number of first database connections established
timesten_connection_established_threshold_exceeded_total	counter	Number of database connection threshold exceeded events
timesten_cs_row_delete_total	counter	Number of table rows deleted by server
timesten_cs_row_insert_total	counter	Number of table rows inserted by server
timesten_cs_server_commit_total	counter	Number of transactions committed by server
timesten_cs_server_disconnected_total	counter	Number of client/server disconnects
timesten_cs_server_execute_alter_total	counter	Number of MERGE statements executed by server
timesten_cs_server_execute_create_total	counter	Number of CREATE statements executed by server
timesten_cs_server_execute_delete_total	counter	Number of DELETE statements executed by server
timesten_cs_server_execute_drop_total	counter	Number of DROP statements executed by server
timesten_cs_server_execute_insert_total	counter	Number of INSERT statements executed by server
timesten_cs_server_execute_select_total	counter	Number of SELECT statements executed by server
timesten_cs_server_execute_update_total	counter	Number of UPDATE statements executed by server
timesten_cs_server_received_byte_total	counter	Number of client/server bytes received by server
timesten_cs_server_rollback_total	counter	Number of transaction rollbacks by server
timesten_cs_server_roundtrip_total	counter	Number of client/server round trips
timesten_cs_server_row_update_total	counter	Number of table rows updated by server
timesten_cs_server_transmitted_byte_total	counter	Number of client/server bytes transmitted by server

Metrics timesten_cursor_close_total to timesten_index_range_rebuild_insert_total

Metric	Type	Description
timesten_cursor_close_total	counter	Number of cursor closes
timesten_cursor_open_total	counter	Number of cursor opens
timesten_database_loaded	gauge	Database loaded
timesten_databases	gauge	Number of databases
timesten_deadlock_total	counter	Number of deadlocks
timesten_disconnect_total	counter	Number of disconnect operations
timesten_first_log_file	gauge	First log file number
timesten_fs_log_read_total	counter	Number of file system log file read operations
timesten_fs_log_write_total	counter	Number of file system log file write operations
timesten_global_delete_executes_total	counter	Number of grid global delete SQL statements executed
timesten_global_executes_total	counter	Number of grid global SQL statements executed
timesten_global_insert_executes_total	counter	Number of grid global insert SQL statements executed
timesten_global_merge_executes_total	counter	Number of grid global merge SQL statements executed
timesten_global_select_executes_total	counter	Number of grid global select SQL statements executed
timesten_global_update_executes_total	counter	Number of grid global update SQL statements executed
timesten_grid_database_instances	gauge	Number of instances defined for database
timesten_grid_database_loaded	gauge	Non-zero if database loaded for instance
timesten_grid_database_opened	gauge	Non-zero if database opened for instance
timesten_grid_database_replica_sets	gauge	Number of replica sets defined for database
timesten_index_hash_delete_total	counter	Number of rows deleted from hash indexes
timesten_index_hash_insert_total	counter	Number of rows inserted into hash indexes
timesten_index_hash_rebuild_insert_total	counter	Number of rows inserted into hash indexes during index rebuild phase of database recovery
timesten_index_hash_row_fetch_repl_total	counter	Number of rows fetched from hash indexes during replication operations
timesten_index_hash_row_fetch_total	counter	Number of rows fetched from hash indexes
timesten_index_hash_scan_repl_total	counter	Number of hash indexes scanned during replication operations (such as insert, update and delete operations on tables)
timesten_index_hash_scan_total	counter	Number of hash indexes scanned
timesten_index_range_delete_total	counter	Number of rows deleted from range indexes
timesten_index_range_insert_total	counter	Number of rows inserted into range indexes
timesten_index_range_rebuild_insert_total	counter	Number of rows inserted into range indexes during index rebuild phase of database recovery

Metrics `timesten_index_range_row_fetch_repl_total` to `timesten_lock_timeout_total`

Metric	Type	Description
<code>timesten_index_range_row_fetch_repl_total</code>	counter	Number of rows fetched from range indexes during replication operations
<code>timesten_index_range_row_fetch_total</code>	counter	Number of rows fetched from range indexes
<code>timesten_index_range_scan_repl_total</code>	counter	Number of range indexes scanned during replication operations (such as insert, update and delete operations on tables)
<code>timesten_index_range_scan_total</code>	counter	Number of range indexes scanned
<code>timesten_index_range_update_total</code>	counter	Number of rows updated on range indexes
<code>timesten_index_rebuild_total</code>	counter	Number of indexes rebuilt
<code>timesten_index_temporary_create_total</code>	counter	Number of temporary indexes created
<code>timesten_index_temporary_row_fetch_repl_total</code>	counter	Number of rows fetched from temporary indexes during replication operations
<code>timesten_index_temporary_row_fetch_total</code>	counter	Number of rows fetched from temporary indexes
<code>timesten_index_temporary_scan_repl_total</code>	counter	Number of temporary indexes scanned during replication operations
<code>timesten_index_temporary_scan_total</code>	counter	Number of temporary indexes scanned
<code>timesten_instance_status</code>	gauge	TimesTen instance status
<code>timesten_local_delete_executes_total</code>	counter	Number of grid local delete SQL statements executed
<code>timesten_local_executes_total</code>	counter	Number of grid local SQL statements executed
<code>timesten_local_insert_executes_total</code>	counter	Number of grid local insert SQL statements executed
<code>timesten_local_merge_executes_total</code>	counter	Number of grid local merge SQL statements executed
<code>timesten_local_select_executes_total</code>	counter	Number of grid local select SQL statements executed
<code>timesten_local_update_executes_total</code>	counter	Number of grid local update SQL statements executed
<code>timesten_lock_acquire_dml_total</code>	counter	Number of locks acquired for DML activity
<code>timesten_lock_acquire_table_scan_total</code>	counter	Number of locks acquired for table scans
<code>timesten_lock_grant_immediate_total</code>	counter	Number of locks granted immediately
<code>timesten_lock_grant_wait_total</code>	counter	Number of locks granted that required waiting
<code>timesten_lock_timeout_total</code>	counter	Number of lock timeouts

Metrics `timesten_log_buffer_insert_byte` to `timesten_remote_update_executes_total`

Metric	Type	Description
<code>timesten_log_buffer_insert_byte</code>	counter	Number of bytes inserted into the log buffer

Metric	Type	Description
<code>timesten_log_buffer_insert_total</code>	counter	Number of log records inserted into the log buffer
<code>timesten_log_buffer_wait_total</code>	counter	Number of log buffer waits
<code>timesten_log_bytes_to_log_buffer_total</code>	counter	Number of bytes written to the log buffer
<code>timesten_log_commit_file_read_total</code>	counter	Number of file system reads from the log for commit processing
<code>timesten_log_commit_read_byte_total</code>	counter	Number of bytes read from the log for commit processing
<code>timesten_log_file_earliest</code>	gauge	Earliest log file that currently exists in the database
<code>timesten_log_file_generated_total</code>	counter	Number of log files generated so far
<code>timesten_log_file_latest</code>	gauge	Most recent log file present
<code>timesten_log_file_read_total</code>	counter	Number of file system reads
<code>timesten_log_file_write_total</code>	counter	Number of file system writes
<code>timesten_log_flush_frame_total</code>	counter	Number of frames in the most recent log flush
<code>timesten_log_force_total</code>	counter	Number of times the log is synched to disk
<code>timesten_log_recovery_read_byte_total</code>	counter	Number of log bytes read during database recovery
<code>timesten_logdir_avail_bytes</code>	gauge	Available bytes on file system holding log files
<code>timesten_logdir_size_bytes</code>	gauge	Total bytes on file system holding log files
<code>timesten_merge_join_total</code>	counter	Number of merge joins done
<code>timesten_metric_collection_seconds</code>	gauge	Time in seconds to collect metrics
<code>timesten_nested_loop_join_total</code>	counter	Number of nested loop joins done
<code>timesten_perm_allocated_bytes</code>	gauge	Bytes of permanent space allocated
<code>timesten_perm_in_use_bytes</code>	gauge	Bytes of permanent space used
<code>timesten_perm_in_use_high_water_bytes</code>	gauge	Maximum bytes of permanent space used
<code>timesten_rate_limited</code>	gauge	State of rate limiting of the exporter
<code>timesten_remote_delete_executes_total</code>	counter	Number of grid remote delete SQL statements executed
<code>timesten_remote_executes_total</code>	counter	Number of grid remote SQL statements executed
<code>timesten_remote_insert_executes_total</code>	counter	Number of grid remote insert SQL statements executed
<code>timesten_remote_merge_executes_total</code>	counter	Number of grid remote merge SQL statements executed
<code>timesten_remote_select_executes_total</code>	counter	Number of grid remote select SQL statements executed
<code>timesten_remote_update_executes_total</code>	counter	Number of grid remote update SQL statements executed

Metrics timesten_repl_log_out_busy_total to timesten_repl_transmitter_catch_lag_total

Metric	Type	Description
timesten_repl_log_out_busy_total	counter	Switched out of track because strand was full or busy
timesten_repl_log_track_arbitration_total	counter	Number of arbitration due to out of order log record
timesten_repl_log_track_wait_total	counter	Number of flusher waits triggered by the track
timesten_repl_log_trackswitch_in_busy_total	counter	Switched into track because strand was full or busy
timesten_repl_log_trackswitch_in_timeout_total	counter	Switched into track because of timeout
timesten_repl_log_trackswitch_out_timeout_total	counter	Switched out of track because of timeout
timesten_repl_log_wait_sleeps_total	counter	Number of times transmitter sleeps waiting for log
timesten_repl_log_waits_total	counter	Number of times transmitter waits for log
timesten_repl_receiver_bulk_insert_applied_to_total	counter	Number of sbLRPtnTuplInsArr applied by the receiver
timesten_repl_receiver_commit_time	counter	Wait time on the receiver
timesten_repl_receiver_commit_total	counter	Number of batches for the receiver
timesten_repl_receiver_longest_batch	counter	Longest time to run a batch
timesten_repl_receiver_longest_running	gauge	Time between begin and commit execution on the receiver
timesten_repl_receiver_max_spin_per_txn	gauge	Max number of spins a transaction faces in the receiver
timesten_repl_receiver_permanent_error_total	counter	Permanent errors raised by receiver
timesten_repl_receiver_serial_modes	counter	Number of times parallel replication had to operate in serial mode
timesten_repl_receiver_transient_error_total	counter	Transient errors raised by receiver
timesten_repl_receiver_txn_begin_waits_total	counter	Number of times the receivers waits for begin dependency
timesten_repl_receiver_txn_commit_no_wait_total	counter	Number transactions for the receiver
timesten_repl_receiver_txn_commit_total	counter	Number of transactions for the receiver
timesten_repl_receiver_txn_commit_wait_spin_total	counter	Number of spins the receivers does when it waits for commit dependency
timesten_repl_receiver_txn_commit_waits_total	counter	Number of times the receivers waits for commit dependency
timesten_repl_receiver_txn_opt_time	counter	Operation time on the receiver
timesten_repl_receiver_wait_backoff_total	counter	The number of times that the receiver has been told to exit rxWaitCTN because it holds a lock that another thread needs
timesten_repl_receiver_wait_stall_2_to_5_total	counter	Receiver begin waits between (approx) 2 and 5 seconds
timesten_repl_receiver_wait_stall_gt_5_total	counter	Receiver begin waits over (approx) 5 seconds

Metric	Type	Description
timesten_repl_receiver_wait_stall_lt_2_total	counter	Receiver begin waits < (approx) 2 seconds
timesten_repl_receiver_wakeup_total	counter	Wakeups by the deadlock detector
timesten_repl_transmitter_bulk_insert_sent_to tal	counter	Number of sLRPtnTupInsArr records sent by the transmitter
timesten_repl_transmitter_catch_lag_total	counter	An estimate of the number of seconds it would take for this transmitters sendLSN to reach the current writeLSN

Metrics timesten_repl_transmitter_drift_request_total to timesten_table_row_insert_total

Metric	Type	Description
timesten_repl_transmitter_drift_request_total	counter	The number of times this transmitter did a drift request
timesten_repl_transmitter_longest_batch	gauge	Longest time between first read and send for batch
timesten_repl_transmitter_longest_running	gauge	Longest Time between reading begin and commit for a transaction
timesten_repl_transmitter_max_actions	gauge	Largest number of actions in a transaction
timesten_repl_transmitter_max_adapted_timeout	gauge	Largest adapted timeout
timesten_repl_transmitter_send_lag	gauge	An estimate of the seconds behind the most advanced transmitter sendLSN for the relevant peer
timesten_replication_conflict_total	counter	Number of replication conflicts
timesten_replication_hold_log_file	gauge	Log file number of replication hold
timesten_replication_hold_log_offset	gauge	Log file offset of replication hold
timesten_replication_peer_connection	gauge	Number of replication peer connections
timesten_replication_peer_retry_total	counter	Replication peer retries
timesten_replication_transaction_total	counter	Number of replication transactions
timesten_required_recovery	gauge	If non-zero, recovery performed on start-up
timesten_sort_total	counter	Number of sorts done
timesten_stmt_execute_alter_total	counter	Number of ALTER statements executed
timesten_stmt_execute_create_total	counter	Number of CREATE statements executed
timesten_stmt_execute_delete_total	counter	Number of DELETE statements executed
timesten_stmt_execute_drop_total	counter	Number of DROP statements executed
timesten_stmt_execute_insert_total	counter	Number of INSERT statements executed
timesten_stmt_execute_merge_total	counter	Number of MERGE statements executed
timesten_stmt_execute_select_total	counter	Number of SELECT statements executed
timesten_stmt_execute_total	counter	Number of SQL statements executed
timesten_stmt_execute_update_total	counter	Number of UPDATE statements executed
timesten_stmt_prepare_command_cache_miss_tota l	counter	Number of command cache misses during statement prepare
timesten_stmt_prepare_total	counter	Number of statement prepares
timesten_stmt_reprepare_automatic_total	counter	Number of automatic statement reprepares

Metric	Type	Description
<code>timesten_stmt_reprepare_total</code>	counter	Number of statement reprepares, including forced and automatic
<code>timesten_swt_byte_total</code>	counter	Number of bytes sent to the Oracle database during SWT cache group operations
<code>timesten_table_row_delete_total</code>	counter	Number of table rows deleted
<code>timesten_table_row_insert_total</code>	counter	Number of table rows inserted

Metrics `timesten_table_row_read_total` to `timesten_txn_voted_no_total`

Metric	Type	Description
<code>timesten_table_row_read_total</code>	counter	Number of table rows read
<code>timesten_table_row_update_total</code>	counter	Number of table rows updated
<code>timesten_table_scan_total</code>	counter	Number of full table scans
<code>timesten_temp_allocated_bytes</code>	gauge	Bytes of temporary space allocated
<code>timesten_temp_in_use_bytes</code>	gauge	Bytes of temporary space in use
<code>timesten_terrors_file_bytes</code>	counter	Number of bytes in <code>terrors.log</code>
<code>timesten_tmesg_file_bytes</code>	counter	Number of bytes in <code>tmesg.log</code>
<code>timesten_tunnel_enqueued_requests_total</code>	counter	Number of requests enqueued for tunneling
<code>timesten_tunnel_false_unresponsive_total</code>	counter	Number of false positives for hangs detected
<code>timesten_tunnel_processed_requests_total</code>	counter	Number of requests processed by the tunnels
<code>timesten_tunnel_timedout_requests_total</code>	counter	Number of requests timed out while tunneling
<code>timesten_tunnel_unresponsive_total</code>	counter	Number of hangs detected for the tunnels
<code>timesten_txn_commit_durable_total</code>	counter	Number of durable transaction commits
<code>timesten_txn_commit_internal_durable_total</code>	counter	Number of durable internal transactions
<code>timesten_txn_commit_internal_replication_total</code>	counter	Number of replication-initiated transaction commits
<code>timesten_txn_commit_internal_total</code>	counter	Number of internal transactions
<code>timesten_txn_commit_internal_xla_total</code>	counter	Number of XLA-initiated transaction commits
<code>timesten_txn_commit_nondurable_total</code>	counter	Number of nondurable transaction commits
<code>timesten_txn_commit_replicated_durable_total</code>	counter	Number of durable replicated transaction commits
<code>timesten_txn_commit_replicated_nondurable_total</code>	counter	Number of nondurable replicated transaction commits
<code>timesten_txn_commit_total</code>	counter	Number of transactions committed
<code>timesten_txn_epoch_total</code>	counter	Number of committed epoch transactions
<code>timesten_txn_forget_alloc_total</code>	counter	Number of transaction forget entries allocated
<code>timesten_txn_forget_used_total</code>	counter	Number of transaction forget entries used
<code>timesten_txn_forget_used_total</code>	counter	Number of transaction forget entries used
<code>timesten_txn_initiated_tm_total</code>	counter	Number of grid transactions initiated as TM
<code>timesten_txn_participated_remote_total</code>	counter	Number of grid transactions in which participated as remote
<code>timesten_txn_rollback_total</code>	counter	Number of transaction rollbacks

Metric	Type	Description
timesten_txn_srso_total	counter	Number of transactions resolved on a single replica set
timesten_txn_tcc_total	counter	Number of transactions resolved by TCC
timesten_txn_voted_no_total	counter	Number of transactions this element has voted NOT commit

Tools for Troubleshooting TimesTen

The following sections in this chapter describe how to use the TimesTen utilities and other tools that are used to diagnose problems with the TimesTen database:

- [Using the ttCapture Utility](#)
- [Using the ttIsql Utility](#)
- [Using the ttStatus Utility](#)
- [Using the Logs Generated by the TimesTen Daemon](#)
- [Using the ttTraceMon Utility](#)
- [Using the ttXactAdmin Utility](#)
- [About Using SNMP Traps for Diagnostics](#)
- [Using ODBC Tracing](#)
- [Monitoring the TimesTen System Tables](#)
- [Using the Query Optimizer](#)

Using the ttCapture Utility

The `ttCapture` utility captures information about the configuration and state of TimesTen. The `ttCapture` utility generates a file named `ttcapture.date.time.log`. By default, the file is written to the directory from which you invoke the `ttCapture` utility. Use the `ttCapture -dest` option to direct the output file to be written to another directory.

If you run `ttCapture` again, it writes the information to a new file.

When you experience a problem with a TimesTen database, run `ttCapture` with the `DSN` option for the database as soon as possible, either when you are encountering the problem or immediately afterward.

 **Note**

Always double-quote directory and file names in case there are spaces in the names.

When you contact TimesTen Customer Support, upload the `ttcapture.date.number.log` file to the Service Request.

See `ttCapture` in the *Oracle TimesTen In-Memory Database Reference* for information about additional options.

Using the ttIsql Utility

The `ttIsql` utility enables you to interactively execute SQL statements and report status information on your TimesTen database.

All TimesTen SQL operations can be executed from a ttIsql Command> prompt.

This example illustrates how to use the ttIsql utility.

```
% ttIsql database1
Copyright (c) 2019, 2021, Oracle and/or its affiliates. All rights reserved.
Type ? or "help" for help, type "exit" to quit ttIsql.

connect "DSN=database1";
Connection successful:
DSN=database1;UID=testuser;DataStore=/scratch/testuser/database1;
DatabaseCharacterSet=AL32UTF8;ConnectionCharacterSet=AL32UTF8;PermSize=128;
(Default setting AutoCommit=1)
```

You can then execute SQL statements or ttIsql commands at the Command> prompt. You can also call TimesTen built-in procedures. One useful procedure is ttConfiguration. When you call ttConfiguration, TimesTen displays the values of the connection attributes for your current database connection. For more information, see ttConfiguration in the *Oracle TimesTen In-Memory Database Reference*.

Using the ttIsql Utility in the *Oracle TimesTen In-Memory Database Operations Guide* describes how to use the most common ttIsql commands. The following ttIsql commands are commonly used when troubleshooting:

- dssize prints database size information.

See Displaying Database Structures Using ttIsql in the *Oracle TimesTen In-Memory Database Operations Guide*.

- showplan prints the optimizer execution plans for selects, updates, and deletes in this transaction.

See Using ttIsql to View and Change Query Optimizer Plans in the *Oracle TimesTen In-Memory Database Operations Guide*.

- timing prints query timing.

See Timing ODBC Function Calls in the *Oracle TimesTen In-Memory Database Operations Guide*.

- optprofile prints the current optimizer flag settings and join order.

See Using ttIsql to View and Change Query Optimizer Plans in the *Oracle TimesTen In-Memory Database Operations Guide*.

For the full list of ttIsql features, see ttIsql in the *Oracle TimesTen In-Memory Database Reference*.

Using the ttStatus Utility

Use the ttStatus utility to check the status of the TimesTen daemon and the state of all TimesTen connections.

In this example, the output from ttStatus indicates that no TimesTen daemon is running. If the daemon has stopped unexpectedly, see [No Response from TimesTen Daemon or Subdaemon](#) for troubleshooting information.

```
% ttStatus
ttStatus: Could not connect to the TimesTen daemon.
If the TimesTen daemon is not running, please start it
by running "ttDaemonAdmin -start".
```

In this example, the output from `ttStatus` indicates that the TimesTen daemon is running. There is one database named `database1`.

The first line indicates that the TimesTen daemon is running as process 26251 on port 6624 for the `instance1` TimesTen instance. The second line indicates the TimesTen Server is running as process 26258 on port 6625.

There are currently 14 connections to the database: one user and 13 subdaemon connections.

The restart policies for the cache agent and the replication agent in the database are set to manual.

```
% ttStatus
TimesTen status report as of Tue Mar 19 10:16:07 2021
```

```
Daemon pid 26251 port 6624 instance instance1
TimesTen server pid 26258 started on port 6625
-----
Data store /scratch/testuser/database1
Daemon pid 26251 port 6624 instance instance1
TimesTen server pid 26258 started on port 6625
There are 14 connections to the data store
Shared Memory KEY 0x021001b9 ID 10813443
PL/SQL Memory KEY 0x031001b9 ID 10846212 Address 0x5000000000
```

Type	PID	Context	Connection Name	ConnID
Process	26291	0x00007f5975c07010	database1	1
Subdaemon	26256	0x00007fe2000008c0	Deadlock Detector	2043
Subdaemon	26256	0x00007fe2040008c0	Checkpoint	2041
Subdaemon	26256	0x00007fe2080008c0	Monitor	2045
Subdaemon	26256	0x00007fe2280008c0	Flusher	2044
Subdaemon	26256	0x00007fe230eb1010	Garbage Collector	2037
Subdaemon	26256	0x00007fe2312ce010	XactId Rollback	2038
Subdaemon	26256	0x00007fe2317ec010	IndexGC	2035
Subdaemon	26256	0x00007fe231d0a010	HistGC	2039
Subdaemon	26256	0x00007fe232228010	AsyncMV	2040
Subdaemon	26256	0x00007fe232746010	Log Marker	2036
Subdaemon	26256	0x00007fe232c64010	Aging	2042
Subdaemon	26256	0x00007fe238170010	Rollback	2046
Subdaemon	26256	0x00007fe23858d010	Manager	2047

```
Replication policy : Manual
Cache Agent policy : Manual
PL/SQL enabled.
```

```
Accessible by group g900
End of report
```

This example shows a connection to an old instance of a database. This can occur when a database is invalidated, but some users have not disconnected from the invalidated copy of the database still in memory. After all users disconnect, the memory is automatically freed.

```
% ttStatus
TimesTen status report as of Mon Aug 6 22:03:04 2021
Daemon pid 5088 port 17000 instance instance1
TimesTen server pid 4344 started on port 17002
-----
Data store c:\temp\sample
There are no connections to the data store
Obsolete or not yet active connection(s):
Process 4696 context 0xd08800 name sample connid 1, obsolete connection, shmKey
'Global\DB145b94c6f.3.SHM.4'
```

```
Replication policy : Manual
Cache agent policy : Manual
-----
End of report
```

Using the Logs Generated by the TimesTen Daemon

The TimesTen daemon and its subdaemons and agents generate 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 be aware of, including actions you may need to take.
- A support log containing everything in the user error log plus information for use by TimesTen Customer Support.

See Error, Warning, and Informational Messages in the *Oracle TimesTen In-Memory Database Operations Guide* for information about configuring the logs, including their location and size.

All message categories are enabled by default, but you can also disable or explicitly enable the following categories of user error log messages for all databases or specified databases through options of the `ttDaemonLog` utility:

- ALL (default): For all messages
- CACHE: For messages from the cache agent, designated by `CAC`
- DAEMON: For messages from the main daemon and subdaemons
- DAEMONDBG: For additional information from the main daemon and subdaemons.
- REPLICATION: For messages from the replication agent, designated by `REP`
- SERVER: For messages from TimesTen Server.

See `ttDaemonLog` in *Oracle TimesTen In-Memory Database Reference* for information about `ttDaemonLog` options.

Using the ttTraceMon Utility

Use the `ttTraceMon` utility to log various trace information on a number of TimesTen components. Each TimesTen component can be traced at different levels of detail. You can list all of the traceable TimesTen components and their current tracing level by specifying `ttTraceMon` with the `show` subcommand. See `ttTraceMon` in the *Oracle TimesTen In-Memory Database Reference* for a list of options.

TimesTen tracing severely impacts application performance and consumes a great deal of file system space if trace output is directed to a file. In addition, when using AWT cache groups, you must restart the replication agent when trying to trace the `ORACON` component with `ttTraceMon`. Use the `ttTraceMon` utility only when diagnosing problems. When you are finished, reset tracing to the default values.

This example shows that the levels for most tracing components are set to level 0 (off) for the `databasel` database. Both the `ERR` and `DEADLOCK` components are set to 1 for tracing by default. See [ERR Tracing](#) for information.

```
% ttTraceMon -e show databasel
AGING      ... 0
```

API	... 0
ASYNCMV	... 0
AUTOREFRESH	... 0
BULKLOAD	... 0
BLK	... 0
CG	... 0
CKPT	... 0
DBG0	... 0
DBG1	... 0
DEADLOCK	... 1
DIST	... 0
EE	... 0
ERR	... 1
EXTTBLLOAD	... 0
FLOW	... 0
GRID	... 0
GRIDAS	... 0
HEAP	... 0
INTERRUPT	... 0
IX	... 0
IXGC	... 0
IXHEAP	... 0
LATCH	... 0
LBCU	... 0
LOB	... 0
LOCK	... 0
LOG	... 0
LOGF	... 0
MEM	... 0
ODBC	... 0
OPT	... 0
ORACON	... 0
PLOAD	... 0
PREP	... 0
PT	... 0
REPL	... 0
SM	... 0
SQL	... 0
TEST	... 0
TRACE	... 1
XA	... 0
XACT	... 0
XLA	... 0

The output for most TimesTen components is only of interest to TimesTen Customer Support. However, the output for the SQL, API, LOCK, ERR, AGING and AUTOREFRESH components may be useful to you when you are troubleshooting application problems.

The rest of this section includes the following topics:

- [Starting a Trace and Reading a Trace Buffer](#)
- [SQL Tracing](#)
- [API Tracing](#)
- [DEADLOCK Tracing](#)
- [LOCK Tracing](#)
- [ERR Tracing](#)
- [AGING Tracing](#)
- [AUTOREFRESH Tracing](#)

Starting a Trace and Reading a Trace Buffer

Start a new trace by specifying `ttTraceMon datastore`. For example, to start a trace on the `database1` database, enter:

```
% ttTraceMon database1
Trace monitor; empty line to exit
Trace >
```

At the `Trace` prompt, specify the type of trace and its level. For example, to start tracing the `SQL` component at level 3, enter:

```
Trace > level sql 3
```

At this point you can run your application and the TimesTen trace information is written to a trace buffer. There are two ways to read the contents of the trace buffer:

- From the `Trace` prompt, use the `outfile` command to direct the trace buffer data to a file. You must do this before running your application. When writing tracing information to a file, new trace information is concatenated to the existing contents of the file.
- From the `Trace` prompt, use the `dump` command to display the trace buffer data to your screen.

Note

The contents of the trace buffer accumulate with each new trace. To clear the trace buffer, use the `flush` command from a `ttTraceMon` prompt. Clear the buffered trace records for a specific component by specifying the component with the `flush` command.

Each record from the trace buffer has the following format:

```
timestamp    sequence    component    level    connection    processid
               operation
```

The fields in the records are defined as follows:

- `timestamp` is the time at which the operation was executed.
- `sequence` is the incremental number that identifies the trace line.
- `component` is the TimesTen component being traced (such as `SQL`, `API`, `LOCK`, or `ERR`).
- `level` is the trace level associated with the trace line. The range of trace levels differs by component, but for all components, the lowest trace level generates the least verbose output and highest trace level generates the most verbose output. For example, if you are tracing `SQL` at level 4, your output includes lines for levels 2 (prepare), 3 (execute), and 4 (open, close, fetch).

Note

Trace levels for some components are not a continuous range of numbers. If you enter a number that does not correspond to a supported level for a component, tracing occurs at the highest supported level that is less than the number you entered. For example, if tracing levels for a component are 1, 2, 3, 4, and 6 and you enter 5, tracing events for level 1, 2, 3, and 4 are generated.

- *connection* is the internal connection ID identifying the connection that generated the trace. This number corresponds to the ConnID shown in the ttStatus output. The connection ID is also used as the first element of the transaction ID.
- *processid* is the operating system process ID for the process that generated the trace.
- *operation* is the operation that occurred (such as SQL statement, API operation, or error message).

For example, a line from the trace buffer after a SQL trace at level 3 might look like this:

```
10:39:50.231 5281 SQL 2L 1C 3914P Preparing: select cust_num from customer
```

SQL Tracing

Using ttTraceMon with the SQL component provides information about the SQL being prepared or executed by the TimesTen engine. [Table 2-1](#) describes the levels for SQL tracing. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-1 SQL tracing levels

Level	Output
2	SQL commands being prepared
3	+ SQL commands being executed
4	+ The effect of command pooling (prepares not being done because the prepared command already exists in the pool), the need for reprepares (for example, because an index was created), and commands being destroyed At this level, ttTraceMon also shows when a query command is being opened, fetched, and closed.
5	+ Some internal data, such as command numbers, which are not generally useful for customer-level debugging

Note

TimesTen recommends tracing SQL at level 3 or 4. SQL tracing does not show any information about the optimizer. Optimizer tracing is managed by a separate component (OPT) at level 4 only, and is not designed for customer use.

In this example, execute ttTraceMon to do a SQL trace at level 4 on the database1 database. Direct the output from the SQL trace to the SQLtrace.txt file. Then flush the buffer so that the trace does not report past SQL statements.

```
% ttTraceMon database1
Trace monitor; empty line to exit
Trace > outfile SQLtrace.txt
Trace > level sql 4
Trace > flush
```

At this point, execute an application that includes the following SQL statement:

```
SELECT * FROM departments WHERE department_id = 10;
```

The trace information is written to the SQLtrace.txt file:

```
12:19:36.582      269 SQL      2L      3C  29570P Preparing: select * from
departments where department_id = 10
12:19:36.583      270 SQL      4L      3C  29570P sbSqlCmdCompile ()(E): (Found
already compiled version: refCount:01, bucket:28) cmdType:100, cmdNum:1000146.
12:19:36.583      271 SQL      4L      3C  29570P Opening: select * from departments
where department_id = 10;
12:19:36.583      272 SQL      4L      3C  29570P Fetching: select * from
departments where department_id = 10;
12:19:36.583      273 SQL      4L      3C  29570P Closing: select * from departments
where department_id = 10;
5 records dumped
```

When the application has completed, turn off SQL tracing and exit ttTraceMon.

```
Trace > level sql 0
Trace > {press ENTER - blank line}
```

API Tracing

API traces are generated for database operations such as connecting to a database, changing a connection attribute, and committing a transaction. [Table 2-2](#) describes the levels for API tracing. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-2 API tracing levels

Level	Output
1	All rollback attempts by the subdaemon This occurs if an application exits abruptly and the subdaemon recovers its connection.
2	+ Some low-on-space conditions
3	+ Create, connect, disconnect, checkpoint, backup, and compact operations for the database, as well as commit and rollback for each connection, and a few other operations
4	+ Most other operations conducted at TimesTen's internal API level It does not show numerous operations on the storage manager and indexes that are done internally.

 **Note**

TimesTen recommends tracing at level 3.

In this example, execute ttTraceMon to do a API trace at level 3 on the `databae1` database. The output from the API trace is written to the `APItrace.txt` file. Before saving the API trace to the buffer, use the `flush` command to empty the buffer.

```
% ttTraceMon databae1
Trace monitor; empty line to exit
Trace> outfile APItrace.txt
Trace> level api 3
Trace > flush
```

At this point, execute the application. When the application has completed, turn off API tracing and exit ttTraceMon:

```
Trace > level api 0
Trace > {press ENTER - blank line}
```

The contents of `APItrace.txt` are similar to the sample output shown below. The output shows connection to the database, setting the connection character set, setting the isolation level, and committing a transaction.

```
11:54:26.796 1016 API      3L    2C  4848P sb_dbConnect()(X)
11:54:26.796 1017 API      3L    2C  4848P sb_dbConnCharsetSet()(E)
11:54:26.796 1018 API      3L    2C  4848P sb_dbConnSetIsoLevel()(E)
11:54:39.795 1019 API      3L    2C  4848P sb_dbConnSetIsoLevel()(E)
11:54:45.253 1020 API      3L    2C  4848P sb_xactCommitQ()(E)
```

DEADLOCK Tracing

Use the DEADLOCK component to trace the occurrences of deadlocks for all applications.

[Table 2-3](#) describes the DEADLOCK tracing levels. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-3 DEADLOCK tracing levels

Level	Output
1	Deadlock cycles, logged as they are discovered
4, 6	+ Detailed information about how the deadlock is detected

In this example, execute ttTraceMon to do a DEADLOCK trace at level 1, which is the default, on `myDSN` database.

Make two connections to `myDSN`. For the first connection, autocommit is on. Create table `test` and insert two rows. Then, set autocommit off and update the `x1=1` row of table `test`. Because autocommit is off, the row is not inserted into the table until the commit. A lock is held until the transaction is committed or rolled back.

```
Command> create table test (x1 int unique, y1 int);
Command> insert into test values (1,1);
1 row inserted.
Command> insert into test values (2,2);
1 row inserted.
Command> autocommit 0;
Command> update test set y1=y1 where x1=1;
1 row updated.
```

For the second connection to `myDSN`, autocommit is set to off. Update the `x1=2` row of table `test`.

```
Command> autocommit 0;
Command> update test set y1=y1 where x1=2;1 row updated.
```

Now, create a deadlock situation by executing update statements in both connections for rows that are locked by each other. The first connection executes an update against the row where x1=2.

```
Command> update test set y1=y1 where x1=2;
  6003: Lock request denied because of time-out
Details: Tran 2.1 (pid 32750) wants Un lock on rowid BMUFVUAAAaAAAAETk,
table ME.TEST. But tran 3.2 (pid 32731) has it in Xn (request was Xn).
Holder SQL (update t1 set y1=y1 where x1=2)
The command failed.
```

The second connection executes an update against the row where x1=1.

```
Command> update test set y1=y1 where x1=1;
  6002: Lock request denied because of deadlock
Details: Tran 3.2 (pid 32731) wants Un lock on rowid BMUFVUAAAaAAAADzk,
table ME.TEST. But tran 2.1 (pid 32750) has it in Xn (request was Xn).
Holder SQL (update t1 set y1=y1 where x1=1)
The command failed.
```

Use the flush command to empty the buffer.

```
% ttTraceMon myDSN
Trace monitor; empty line to exit
Trace> flush
```

The trace buffer contains the following information showing all level 1 deadlock traces, as evidenced by '1L'.

```
Trace> dump
09:50:26.444      13 DEADLOCK 1L 2036C  3484P edge 1: xid 3.2, cid 3,
<Row BMUFVUAAAaAAAADzk,0x8c5
74(574836)> 0 cnt=1 , Tbl 'T1', SQL='update t1 set y1=y1 where x1=1'
09:50:26.455      14 DEADLOCK 1L 2036C  3484P edge 0: xid 2.1, cid 2,
<Row BMUFVUAAAaAAAETk,0x8c5
74(574836)> 0 cnt=1 , Tbl 'T1', SQL='update t1 set y1=y1 where x1=2'
09:50:26.455      15 DEADLOCK 1L 2036C  3484P Victim: xcb:3.2,
SQL='update t1 set y1=y1 where x1=1'
```

If you want more information, set DEADLOCK tracing to a higher value. For example, the following sets DEADLOCK tracing to level 4 in ttTraceMon:

```
Trace > level deadlock 4
```

LOCK Tracing

Use the LOCK component to trace the locking behavior of your application to detect trouble with deadlocks or lock waits. LOCK tracing generates a great deal of volume, so it is important to choose the appropriate level of tracing. Level 3, for example, begins to generate a large number of traces, as traces are written for fairly common events. Also, the traces themselves may be hard to read in places. If you cannot discern enough information, contact TimesTen Customer Support for more information.

[Table 2-4](#) describes the LOCK tracing levels. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-4 LOCK tracing levels

Level	Output
1	Deadlock cycles as they are discovered
2	+ Failures to grant locks for any reason
3	+ Lock waits (which may or may not be granted)
4	+ All lock grants/releases, some routine calls, and the mechanism of the deadlock detector
6	+ Each step in cycle traversal

In this example, execute ttTraceMon to do a `LOCK` trace at level 4 on `myDSN` database.

Make two connections to `myDSN`. For the first connection, set autocommit on. Create table `test` and insert three rows. Create a materialized view using table `test`.

Turn on tracing at level 4 and use the `flush` command to empty the buffer.

```
% ttTraceMon myDSN
Trace monitor; empty line to exit
Trace> level lock 4
Trace> flush
```

For the second connection to `myDSN`, set autocommit off. Insert a row into table `test`. Because autocommit is off, the row is not inserted into the table until the commit. A lock is held until the transaction is committed or rolled back.

If the `dump` command is used to display the contents of the trace buffer, there are no records in the trace buffer:

```
Trace> dump
0 records dumped
```

From the first connection, try to drop the materialized view. The view cannot be dropped because there is a transaction that has not been committed or rolled back:

```
Command> drop materialized view v;
6003: Lock request denied because of time-out
Details: Tran 3.71 (pid 24524) wants Sn lock on table TTUSER.TEST. But tran 1.42
(pid 24263) has it in IXn (request was IXn). Holder SQL (insert into test
values (100););
The command failed.
```

The trace buffer contains the following information:

```
Trace> dump
20:09:04.789 174759 LOCK      3L      3C  24524P ENQ: xcb:00003 <Tbl 0x9b894,0x0>
0+Sn=>SL activity 0 Sn cnt=0; Holder xcb:00001 IXn
20:09:04.789 174760 LOCK      3L      3C  24524P Connection 3 scheduled for sleep
20:09:04.789 174761 LOCK      3L      3C  24524P Connection 3 sleeping
20:09:14.871 174762 LOCK      3L 2047C  24237P Connection 3 timed out
20:09:14.871 174763 LOCK      3L 2047C  24237P Connection 3 woken up
20:09:14.871 174764 LOCK      3L      3C  24524P Connection 3 awake
20:09:14.871 174765 LOCK      2L      3C  24524P ENQ: xcb:00003 <Tbl 0x9b894,0x0>
0+Sn=>TM activity 0 Sn cnt=1; Holder xcb:00001 IXn
7 records dumped
```

When finished with the trace, set `LOCK` tracing back to its default setting (0) and exit `ttTraceMon`:

```
Trace > level lock 0
Trace > {press ENTER - blank line}
```

ERR Tracing

It may be useful to trace the `ERR` component. For example, an `ERR` trace at level 4 shows all of the error messages that are pushed in the TimesTen direct driver (not all errors are output to the user because they are handled internally). `ERR` tracing at level 1 is the default. No output is written for `ERR` tracing at level 2 and 3.

[Table 2-5](#) describes `ERR` tracing levels. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-5 ERR tracing levels

Level	Output
1 (set by default)	Fatal errors
4	+ All other error messages, many of which are handled internally by TimesTen

In this example, execute `ttTraceMon` to do a `ERR` trace at level 4 on `myDSN` database.

First, create a table:

```
Command> CREATE TABLE test (id TT_INTEGER);
```

Next, turn on tracing at level 4. Rather than direct the trace output to a file as in the previous examples, read it directly from the trace buffer. Before saving the `ERR` trace to the buffer, use the `flush` command to empty the buffer.

```
% ttTraceMon myDSN
Trace monitor; empty line to exit
Trace> level err 4
Trace> flush
```

Now execute a SQL script with three errors in it. The errors are:

- Creating a table with the same name as an existing table
- Using incorrect syntax to insert values into the table
- Inserting `CHAR` data into a `TT_INTEGER` column

```
Command> CREATE TABLE test (id TT_INTEGER);
2207: Table TEST already exists
The command failed.
Command> INSERT INTO test VALUES 'abcd';
1001: Syntax error in SQL statement before or at: "'abcd'", character position:
25
insert into test values 'abcd';
      ^^^^^^
The command failed.
Command> INSERT INTO test VALUES ('abcd');
2609: Incompatible types found in expression
The command failed.
```

The trace information is written to the trace buffer. Display it by using the `dump` command.

```

Trace> dump
19:28:40.465 174227 ERR      4L    1C 24263P TT2207: Table TEST already exists
-- file "eeDDL.c", lineno 2930, procedure "sbEeCrDtblEval()"
19:28:51.399 174228 ERR      4L    1C 24263P TT1001: Syntax error in SQL
statement before or at: "'abcd'", character position: 25
insert into test values 'abcd');
                                ^^^^^^
-- file "ptSqlY.y", lineno 6273, procedure "reserved_word_or_syntax_error"
19:29:00.725 174229 ERR      4L    1C 24263P TT2609: Incompatible types found
in expression -- file "saCanon.c", lineno 12618, procedure "sbPtAdjustType()"
3 records dumped

```

Set ERR tracing back to its default setting (1) and exit ttTraceMon:

```

Trace > level err 1
Trace > {press ENTER - blank line}

```

AGING Tracing

Use the ttTraceMon utility to obtain the following information:

- When aging starts and ends
- How many rows have been deleted by the aging subdaemon

See *Implementing an Aging Policy in Your Tables* in the *Oracle TimesTen In-Memory Database Operations Guide*.

[Table 2-6](#) describes the AGING tracing levels. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-6 AGING Tracing Levels

Level	Description
1	Messages about the following events: <ul style="list-style-type: none"> • The aging subdaemon starts least recently used (LRU) or time-based aging. • The aging subdaemon repeats LRU aging because the LRU threshold was not met. • The aging subdaemon ends LRU or time-based aging.
2	+ Messages about the following events <i>for each table</i> : <ul style="list-style-type: none"> • Aging has started. • Aging has ended. The message includes the reason for ending and the total number of rows deleted.
3	+ Detailed report on how many rows were deleted during each aging cycle
4	+ Message every time the aging subdaemon wakes up

In this example, execute ttTraceMon to do an AGING trace on myDSN database. The database contains TTUSER.MYTAB table, which has a time-based aging policy. The table is described as follows:

```
Command> DESCRIBE ttuser.mytab;
```

```

Table TTUSER.MYTAB:
Columns:
  *ID          TT_INTEGER NOT NULL
  TS          TIMESTAMP (6) NOT NULL
Aging use TS lifetime 3 minutes cycle 1 minute on

```

```
1 table found.  
(primary key columns are indicated with *)
```

The table contains the following rows before the aging cycle begins:

```
Command> select * from TTUSER.MYTAB;  
< 1, 2007-03-21 12:54:06.000000 >  
< 3, 2010-03-17 08:00:00.000000 >  
< 4, 2007-03-21 12:59:40.000000 >  
< 5, 2007-03-21 13:00:10.000000 >  
< 6, 2007-03-21 13:01:22.000000 >  
5 rows found.
```

Execute ttTraceMon to do an AGING trace at level 3. Rather than direct the trace output to a file, read it directly from the trace buffer. Before saving the AGING trace to the buffer, use the flush command to empty the buffer.

```
% ttTraceMon myDSN  
Trace monitor; empty line to exit  
Trace> level aging 3  
Trace> flush
```

Display the trace information in the buffer by using the dump command.

```
Trace> dump  
13:16:56.802      1247 AGING      1L 2045C 17373P Entering sbAgingTB(): curTime=78  
13:16:56.803      1248 AGING      2L 2045C 17373P Entering sbAgingOneTable():  
curTime=78, ltblid= 637140  
13:16:56.804      1249 AGING      3L 2045C 17373P curTime=78, 4 deleted, 1  
remaining, tbl = TTUSER.MYTAB  
13:16:56.804      1250 AGING      2L 2045C 17373P Exiting sbAgingOneTable():  
curTime=78, reason = 'no more rows', 4 deleted, 1 remaining, tbl = TTUSER.MYTAB  
13:16:56.804      1251 AGING      1L 2045C 17373P Exiting sbAgingTB(): curTime=78  
5 records dumped
```

Set AGING tracing back to its default setting (0) and exit ttTraceMon:

```
Trace > level aging 0  
Trace > {press ENTER - blank line}
```

AUTOREFRESH Tracing

Use the ttTraceMon utility to obtain information about the progress of autorefresh operations for cache groups with the AUTOREFRESH cache group attribute.

See AUTOREFRESH Cache Group Attribute Overview in the *Oracle TimesTen In-Memory Database Cache Guide*.

[Table 2-7](#) describes AUTOREFRESH tracing levels. Each level with a '+' sign includes the trace information described for that level, plus all levels preceding it.

Table 2-7 AUTOREFRESH Tracing Levels

Level	Description
1	<p>Autorefresh summary for the interval:</p> <ul style="list-style-type: none"> • The time that autorefresh started • Number of autorefreshed rows for the interval • Number of autorefreshed root table rows for interval • Total number of autorefreshed rows since the cache agent started • Total number of autorefreshed rows in the root table since the cache agent started • The time that autorefresh ended <p>Note: Times and information about root table rows are reported for full autorefresh.</p>
2	<p>+ Autorefresh summary at the cache group level:</p> <ul style="list-style-type: none"> • The time that autorefresh started for each cache group • Number of autorefreshed rows for each cache group • Number of autorefreshed root table rows for each cache group • Total number of autorefreshed rows for each cache group since the cache agent started • Total number of autorefreshed rows in the root table for each cache group since the cache agent started • The time that autorefresh ended for each cache group <p>Note: Times and information about root table rows are reported for full autorefresh.</p>
3	<p>+ Autorefresh summary at the table level:</p> <ul style="list-style-type: none"> • The time that autorefresh started • Number of autorefreshed rows • Total number of autorefreshed rows since the cache agent started • The time that autorefresh ended
4	<p>+ Autorefresh details for each table:</p> <ul style="list-style-type: none"> • The time that autorefresh started for each table • The autorefresh query • Query execute time in milliseconds on the Oracle database • Query fetch time in milliseconds on the Oracle database • Query apply time in milliseconds on TimesTen • Query execute time in milliseconds on the Oracle database for child tables • Query fetch time in milliseconds on the Oracle database for child tables • Query apply time in milliseconds on TimesTen for child tables • The time that autorefresh ended for each table • The autorefresh bookmark (logseq) to which autorefresh was completed

In this example, use the ttTraceMon utility to trace autorefresh operations on the `cgDSN` database. When the trace level is set to 1, the displayed information is similar to the output of the `ttCacheAutorefreshStatsGet` built-in procedure.

```
% tttracemon cgDSN
Trace monitor; empty line to exit
Trace> level autorefresh 1
Trace> dump

08:56:57.445 19398 AUTOREFRESH 1L 5C 32246P Autorefresh number 1415 started
for interval 60000
08:56:57.883 19419 AUTOREFRESH 1L 5C 32246P Duration For Interval 60000ms: 420
08:56:57.883 19420 AUTOREFRESH 1L 5C 32246P Num Rows For Interval 60000ms: 0
08:56:57.883 19421 AUTOREFRESH 1L 5C 32246P Num Root Rows For Interval
60000ms: 0
```

```
08:56:57.883 19422 AUTOREFRESH 1L 5C 32246P Cumulative Rows for Interval
60000ms: 11587
08:56:57.883 19423 AUTOREFRESH 1L 5C 32246P Cumulative Root Rows for Interval 60000ms:
1697
08:56:57.883 19424 AUTOREFRESH 1L 5C 32246P Autorefresh number 1415 ended for
interval 60000ms successfully.
7 records dumped
```

Tracing at level 4 produces additional information about autorefresh operation 1415. Information about autorefresh is provided for the testuser.readcache cache group, the testuser.readtab root table and the autorefresh interval.

```
Trace> level autorefresh 4
Trace> dump
```

```
08:56:57.445 19398 AUTOREFRESH 1L 5C 32246P Autorefresh number 1415 started for
interval 60000
08:56:57.471 19399 AUTOREFRESH 2L 5C 32246P Autorefresh started for cachegroup
TESTUSER.READCACHE
08:56:57.471 19400 AUTOREFRESH 3L 5C 32246P Incremental autorefresh started for
table TESTUSER.READTAB
08:56:57.471 19401 AUTOREFRESH 4L 5C 32246P Autorefresh Query: SELECT L."COL_10",
X."COL_20", X.ft$NotDelete, Z.rowid FROM (SELECT DISTINCT "COL_10" FROM
"TESTUSER"."TT_06_454854_L" WHERE logseq >:logseq AND ft_cacheGroup <>
100000000000*1844259679+-299200618) L,(SELECT "TESTUSER"."READTAB"."COL_10",
"TESTUSER"."READTAB"."COL_20", 1 AS ft$NotDelete FROM "TESTUSER"."READTAB",
"TESTUSER"."T1" WHERE "TESTUSER"."READTAB"."COL_10" = "TESTUSER"."T1"."COL_10")
X, "TESTUSER"."READTAB" Z WHERE L ."COL_10" = X."COL_10" (+) AND X."COL_10" =
Z."COL_10" (+), logseq: 7
08:56:57.870 19402 AUTOREFRESH 3L 5C 32246P Duration for table
TESTUSER.READTAB: 70
08:56:57.870 19403 AUTOREFRESH 3L 5C 32246P Num Rows for table
TESTUSER.READTAB: 1
08:56:57.870 19404 AUTOREFRESH 3L 5C 32246P Cumulative rows for table
TESTUSER.READTAB: 1559
08:56:57.870 19405 AUTOREFRESH 4L 5C 32246P Autorefresh Query Execute duration
for table TESTUSER.READTAB: 60
08:56:57.870 19406 AUTOREFRESH 4L 5C 32246P Autorefresh Query Fetch duration for
table TESTUSER.READTAB: 0
08:56:57.870 19407 AUTOREFRESH 4L 5C 32246P Autorefresh Query Apply duration for
table TESTUSER.READTAB: 0
08:56:57.870 19408 AUTOREFRESH 4L 5C 32246P Max logseq applied for table
TESTUSER.READTAB: 8
08:56:57.870 19409 AUTOREFRESH 4L 5C 32246P Autorefresh Query Execute duration
for 7 child(ren) table(s): 32
08:56:57.870 19410 AUTOREFRESH 4L 5C 32246P Autorefresh Query Fetch duration for
7 child(ren) table(s): 0
08:56:57.870 19411 AUTOREFRESH 4L 5C 32246P Autorefresh Query Apply duration for
7 child(ren) table(s): 0
08:56:57.870 19412 AUTOREFRESH 3L 5C 32246P Incremental autorefresh ended for
table TESTUSER.READTAB
08:56:57.872 19413 AUTOREFRESH 2L 5C 32246P Duration For Cache Group
TESTUSER.READCACHE: 1020
08:56:57.872 19414 AUTOREFRESH 2L 5C 32246P Num Rows For Cache Group
TESTUSER.READCACHE: 1
08:56:57.872 19415 AUTOREFRESH 2L 5C 32246P Num Root Rows For Cache Group
TESTUSER.READCACHE: 0
08:56:57.872 19416 AUTOREFRESH 2L 5C 32246P Cumulative Rows for Cache Group
TESTUSER.READCACHE: 11776
08:56:57.872 19417 AUTOREFRESH 2L 5C 32246P Cumulative Root Rows for Cache Group
TESTUSER.READCACHE: 1697
08:56:57.872 19418 AUTOREFRESH 2L 5C 32246P Autorefresh ended for cache group
```

```
TESTUSER.READCACHE
08:56:57.883 19419 AUTOREFRESH 1L 5C 32246P Duration For Interval 60000ms: 420
08:56:57.883 19420 AUTOREFRESH 1L 5C 32246P Num Rows For Interval 60000ms: 0
08:56:57.883 19421 AUTOREFRESH 1L 5C 32246P Num Root Rows For Interval
60000ms: 0
08:56:57.883 19422 AUTOREFRESH 1L 5C 32246P Cumulative Rows for Interval
60000ms: 11587
08:56:57.883 19423 AUTOREFRESH 1L 5C 32246P Cumulative Root Rows for Interval
60000ms: 1697
08:56:57.883 19424 AUTOREFRESH 1L 5C 32246P Autorefresh number 1415 ended for
interval 60000ms successfully.
27 records dumped
```

Set AUTOREFRESH tracing back to its default setting (0) and exit ttTraceMon:

```
Trace > level autorefresh 0
Trace > {press ENTER - blank line}
```

Using the ttXactAdmin Utility

The ttXactAdmin utility displays ownership, status, log and lock information for each outstanding transaction. You can also use it to show all current connections to a database. The ttXactAdmin utility is useful for troubleshooting problems with replication, XLA, and asynchronous writethrough cache groups.

Use ttXactAdmin to diagnose a lock timeout. Consider two connections that are trying to update the same row. The following transaction by Connection 1 is in progress:

```
UPDATE table1 SET c1 = 2 WHERE c1 = 1;
```

Connection 2 attempts to make the following update:

```
UPDATE table1 SET c1 = 3 WHERE c1 = 1;
```

Connection 2 receives the following error:

```
6003: Lock request denied because of time-out
  Details: Tran 2.3 (pid 2880) wants Un lock on rowid 0x00156bbc, table
  TTUSER.TABLE1.
  But tran 1.21 (pid 2564) has it in Xn (request was Xn). Holder SQL
  (update table1 set c1 = 2 where c1 = 1;)
  The command failed.
```

The details of the error indicate that transaction 1.21 has a lock on row 0x00156bbc, the row that transaction 2.3 wants to update. ttXactAdmin displays this information in output that pertains to actions in the entire database:

```
% ttXactAdmin myDSN
2011-03-07 12:57:41.237
c:\datastore\myDSN
TimesTen Release 11.2.2.0.0

Outstanding locks

  PID  Context  TransID  TransStatus  Resource  ResourceID  Mode  Name
  Program File Name: ttIsql

  2564  0xeeb9a8  1.21      Active      Database  0x01312d00  IX
                                         Row          0x00156bbc  Xn  TTUSER.TABLE1
                                         Table        1910868   IXn  TTUSER.TABLE1
```

```
Program File Name: ttIsql

2880 0xeeeb9a8 2.3      Active      Database  0x01312d00  IX
                                Table    1910868    IXn  TTUSER.TABLE1
                                Command  19972120  S

Awaiting locks

PID  Context   TransID Resource ResourceID RMode HolderTransID HMode Name
2880 0xeeeb9a8 2.3      Row      0x00156bbc  Un      1.21          Xn  TTUSER.TABLE1

2 outstanding transactions found
```

See `ttXactAdmin` in the *Oracle TimesTen In-Memory Database Reference*.

About Using SNMP Traps for Diagnostics

Simple Network Management Program (SNMP) is a protocol for network management services. Even though network management software typically uses SNMP to query or control the state of network devices like routers and switches, SNMP can be used to manage almost any resource, such as applications and databases. These devices or resources also can generate asynchronous alerts called traps to inform the management systems of problems. TimesTen supports such SNMP traps.

SNMP traps can use either UDP/IP or TCP/IP as communication protocol. TimesTen supports SNMP version 1, 2, and 3 traps. TimesTen does not send SNMP traps by itself. To generate and receive SNMP traps, you need to install Net-SNMP with the `snmptrap` utility. You also need to enable SNMP trap generation in the `timesten.conf` TimesTen instance configuration file. See *Diagnostics Through SNMP Traps and How TimesTen Sends SNMP Traps* in the *Oracle TimesTen In-Memory Database Error Messages and SNMP Traps* for details.

Using ODBC Tracing

On Linux and UNIX systems, ODBC tracing is available only when using a driver manager. To turn on tracing, set the `Trace` and `TraceFile` attributes.

Monitoring the TimesTen System Tables

Each TimesTen database contains a group of system tables that store metadata about the current state of the database. The system tables are described in *System Tables and Views* in the *Oracle TimesTen In-Memory Database System Tables and Views Reference*.

Note

You can execute `SELECT` statements on a system table, but you cannot execute a statement such as `INSERT`, `UPDATE` or `DELETE` on these tables.

Use the system table `SYS.SYSTEMSTATS` to view database metrics. In addition, the `ttStats` utility enables you to monitor database metrics and to take and compare snapshots of metrics. For more information, see `SYS.SYSTEMSTATS` in the *Oracle TimesTen In-Memory Database System Tables and Views Reference*, `ttStats` in the *Oracle TimesTen In-Memory Database*

Reference, and TT_STATS in the *Oracle TimesTen In-Memory Database PL/SQL Packages Reference*.

You can also view the SYS.MONITOR table, which contains statistics about certain events that have occurred since the first connection to the database. For example, the SYS.MONITOR table contains information about the number of connections to the database; the number of checkpoints taken; the size of the database; and the amount of memory currently in use. Check the contents of the SYS.MONITOR table by executing SELECT statements on the columns or by using the ttIsql monitor command. For an example of how to use the ttIsql monitor command, see Using the ttIsql Utility in the *Oracle TimesTen In-Memory Database Operations Guide*.

Using the Query Optimizer

The query optimizer is an important tool for performance tuning.

For details about using the query optimizer, see:

- The TimesTen Query Optimizer in the *Oracle TimesTen In-Memory Database Operations Guide*
- Using ttIsql to View and Change Query Optimizer Plans in the *Oracle TimesTen In-Memory Database Operations Guide*

If you find that a given query runs more slowly than expected, confirm that the query optimizer has the latest statistics for the tables in your query. See [Update Query Optimizer Statistics](#) for more information. If, after updating your statistics, your query still runs too slowly, it is possible that the TimesTen optimizer is not choosing the optimal query plan to answer that query. Under these circumstances, you can adjust how the optimizer generates a plan by using the ttOpt procedures described in Modifying a SQL Query Execution Plan in the *Oracle TimesTen In-Memory Database Operations Guide*. You can also use statement level optimizer hints to influence the optimizer at the statement level. For more information, see Statement Level Optimizer Hints in the *Oracle TimesTen In-Memory Database SQL Reference*.

Troubleshooting TimesTen Applications and Databases

The following sections provide information to help you diagnose and remedy some problems encountered while using a TimesTen database:

 **Note**

If you are still having problems with your database after following the troubleshooting recommendations in this chapter, contact TimesTen Customer Support.

- [Unable to Start or Stop TimesTen Daemon](#)
- [No Response from TimesTen Daemon or Subdaemon](#)
- [Unable to Create Shared Segment](#)
- [Application Unable to Connect to Database in Direct Mode](#)
- [Troubleshooting Client/Server Problems](#)
- [Application Connects or Disconnects Are Slow](#)
- [Application Is Disconnected Unexpectedly](#)
- [Application Is Slow](#)
- [Application Unresponsive, Appears Hung](#)
- [Application Unable to Find Previously Created Objects](#)
- [Troubleshooting OCI and Pro*C/C++ Applications](#)
- [Running Out of a Resource](#)
- [Duplicate Results from a SELECT Statement](#)
- [Cannot Attach PL/SQL Shared Memory](#)

Unable to Start or Stop TimesTen Daemon

This section describes what to check if you cannot start or stop the TimesTen main daemon.

Possible cause	What to do
Privilege is incorrect.	You need the ADMIN privilege to start or stop the TimesTen daemon. Ensure that you are using the <code>ttDaemonAdmin</code> utility to start the daemon. The output from <code>ttDaemonAdmin</code> shows whether you have the correct privilege.

Possible cause	What to do
Another process is using the TimesTen daemon port.	Use the <code>ttVersion</code> utility to verify what port number the TimesTen daemon is expected to use. Use an operating system command like <code>netstat</code> to check whether another process is listening on the port. If there is a conflict, either change the port number used by the other process.
TimesTen daemon is already running.	Ensure that you are using the <code>ttDaemonAdmin</code> utility to start the daemon. The output from <code>ttDaemonAdmin</code> shows whether the daemon is already running.
There are other problems.	Inspect the user error log produced by the daemon.

No Response from TimesTen Daemon or Subdaemon

The following sections describe what to do if one or more of the TimesTen processes appears to be unavailable:

- [Check the TimesTen User Error Log](#)
- [Extract a Stack Trace from a Core File](#)

Check the TimesTen User Error Log

If you receive an error that indicates the TimesTen subdaemon has stopped and inspect the user error log.

If the TimesTen daemon crashes, it cannot send anything to the user error log, but the subdaemons send a 'main daemon vanished' message to the log before exiting:

```
09:24:13 Err : 4375 -----: Main daemon has vanished
```

Restart the daemon. The next connection to each database causes TimesTen to recover from the checkpoint and transaction log files. See Working with the TimesTen Daemon in the *Oracle TimesTen In-Memory Database Operations Guide*.

Extract a Stack Trace from a Core File

If you experience a crash by one of the TimesTen processes on a Linux or UNIX system and have exhausted all of the diagnostic options, check to see if TimesTen has generated a core file. Use the `ttVersion` utility to find the core file. Look for a line in the output that shows a path for the daemon home directory:

```
TimesTen Release (ttuser:40732)
2011-04-04T17:53:04Z
  Instance admin: ttuser
  Instance home directory:
  /node1/ttuser/ttcur/TTBuild/linux8664_dbg/install
  Daemon home directory:
  /node1/ttuser/ttcur/TTBuild/linux8664_dbg/install/info
```

After locating the core file, attach to the debugger on the system and extract the stack trace from the core file and send the trace results to TimesTen Customer Support.

Unable to Create Shared Segment

You may receive an error that indicates that a shared segment could not be created:

```
4671: TT14000: TimesTen daemon internal error: Error 28 creating shared segment,  
KEY 0x0201f7eb  
4671: -- OS reports too many shared segments in use  
4671: -- Confirm using 'ipcs' and take appropriate action  
4671: 18538 -----: subdaemon process exited
```

Using the Linux or UNIX `ipcs` command may display information like this:

```
----- Shared Memory Segments -----  
key      shmid   owner    perms      bytes    nattch    status  
0x00000000 1098350592 user1    777        10624      2        dest  
0x00000000 1084817409 user1    777        2439680      2        dest  
0x911fc211 1098383362 user2    666        67108864      1  
0x2814afba 170721285  root     666        1048576      1
```

A status of dest means the memory segment is marked to be destroyed. `nattch` shows the number of processes still attached to the memory segment. The `ipcrm` command cannot free the shared memory until the processes detach from the segment or exit. If an application connects to TimesTen and then becomes inactive, nothing can free the shared memory until the user exits or stops the application.

Application Unable to Connect to Database in Direct Mode

This section describes what to check if your application cannot connect to a database in direct mode.

Possible cause	See...
There is a mismatch between the release of TimesTen and database.	Upgrading Your Database for information.
User does not have <code>CREATE SESSION</code> privilege.	Privileges to Connect to Database for information.
File permissions are incorrect.	Check File System Permissions to Access Database for information.
TimesTen daemon or Data Manager service is not running.	Check the TimesTen Daemon Is Running for information.
Incompatible connection attributes or incorrect path name for database set in the DSN.	Check DSN Definition for information.
There is no available shared memory segment or maximum size of shared memory segment is too small.	Manage Semaphores and Shared Memory Segments for information.
There is not enough swap space.	Check Available Swap Space (Virtual Memory) for information.
Number of file descriptors is inadequate.	Increase the Number of Available File Descriptors for information.
There are other possible causes.	Using the Logs Generated by the TimesTen Daemon for information.

Upgrading Your Database

A database is only guaranteed to be accessible by the same minor release of TimesTen that was used to create the database. When you upgrade the TimesTen software and you would like to use the new release to access a database that was previously created, create a database with the new release. Then use the `ttMigrate` utility to copy the tables, indexes, and table data from the old database to the new one.

See Upgrades in TimesTen Classic in the *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide* for details.

Privileges to Connect to Database

The user must have the `CREATE SESSION` privilege to connect to the database. If you do not have access, the administrator must use the `GRANT` statement to grant you the `CREATE SESSION` privilege. See Privileges to Connect to the Database in the *Oracle TimesTen In-Memory Database Security Guide*.

Check File System Permissions to Access Database

A permission denied error is generated if you attempt to connect to a database and you do not have the proper permissions to access the checkpoint or transaction log files or the directory where those files reside. Check the file system permissions on the files located in the directory specified in the `DataStore` attribute in your DSN.

Check the TimesTen Daemon Is Running

If the TimesTen daemon or Data Manager service is not running, an attempt to connect to a database generates TimesTen error 799 "Unable to connect to daemon; check daemon status."

Use the `ttStatus` utility as described in [Check the TimesTen User Error Log](#) to check the status of the TimesTen daemon.

Check DSN Definition

In your DSN description, perform the following:

- [Check DSN Attributes](#)
- [Check Path Name to Database and Transaction Log Directories](#)

Check DSN Attributes

Certain connection options or DSN attribute settings combinations are not compatible. In cases where incompatible settings are used, an error is returned to the application when it attempts to connect to a database.

Check Path Name to Database and Transaction Log Directories

Confirm that you have specified the correct path names in the `DataStore` and `LogDir` attributes in your DSN. Also confirm that the path names are absolute path names, rather than relative. Otherwise, the path name will be relative to the directory where the application was started.

Manage Semaphores and Shared Memory Segments

An error is generated if you attempt to connect to or create a shared database whose size is larger than the maximum size of shared memory segments configured on your system. Also, an error is generated if the system cannot allocate any more shared memory segments.

On Linux and UNIX systems, use commands similar to the following:

- `ipcs -ma` to check if you have other shared memory segments using up memory, such as Oracle Database instances or other instances of TimesTen.
- `ipcrm` to remove a message queue, semaphore set or shared memory segment identifier. Use `ipcrm` to clean up semaphores or shared memory segments after a faulty TimesTen shutdown, instance crash, daemon crash or other application issues that use shared memory segments and semaphores. Use `-m` to remove a shared memory segment. Use `-s` to remove a semaphore.
- `ps -eafl` to see how much memory is being used by running processes.
- `ulimit -a` to see if there are any limits on the maximum amount of memory one process can address, maximum file size, and the maximum number of open files.

If a shared memory segment is available but is too small to hold your database, use the `ttSize` utility to estimate the amount of memory required for your tables and then check the values of the `PermSize` and `TempSize` attributes to verify the amount of memory established for your database. Monitoring `PermSize` and `TempSize` Attributes in the *Oracle TimesTen In-Memory Database Operations Guide* describes guidelines for setting the size of your permanent and temporary memory regions. If the amount of memory established for your database is too large, reset `PermSize` and `TempSize` to smaller values. See [Check the Amount of Memory Allocated to a Database](#) for more information. Another option is to increase the maximum size of the shared memory segment, as described below.

If your application is connected to your database and your database is invalidated because of a system or application failure, the database shared segment is not automatically detached from your application. The next time your application tries to use the database, the application will be disconnected from the database and the shared segment will be detached. If your application remains connected to the invalidated database and does not try to use the database, the application will not disconnect from the database and the shared segment will not be detached. To free memory and swap space, make sure you disconnect or terminate all applications that are connected to the invalidated database.

If a database becomes invalidated because of a system or application failure, and you have disconnected all applications from the database, a subsequent connection recovers the database. If recovery fails because you have run out of database space, then reconnect to the database with a larger `PermSize` and `TempSize` value than the ones that are currently in effect.

See Operating System Prerequisites in the *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide* for more information on how to configure shared memory.

Check Available Swap Space (Virtual Memory)

There must be enough swap space to back up shared memory.

On Linux and UNIX systems, use the `swap` command to check and add virtual memory to your system.

Increase the Number of Available File Descriptors

Each process connected to a TimesTen database keeps at least one operating system file descriptor open. Additional file descriptors may be opened for a connection when the connection commits or rolls back a transaction. If you receive an error that all file descriptors are in use when attempting to connect to a database, increase the allowable number of file descriptors. See your operating system documentation for limits on file descriptors and information about changing the number of file descriptors. Also see Limits on Number of Open Files in the *Oracle TimesTen In-Memory Database Reference*.

Troubleshooting Client/Server Problems

This section includes the following topics:

- [Cannot Connect to the TimesTen Server](#)
- [TimesTen Server Failed](#)
- [Cannot Find Server DSN](#)
- [TimesTen Server Failed to Load Driver](#)
- [Application Times Out When Accessing TimesTen Server](#)
- [TimesTen Client Loses Connection with TimesTen Server](#)
- [Failed to Attach to Shared Memory Segment for IPC](#)
- [Thread Stack Overflow When Using Multiple Client Connections](#)
- [Out of Space When DSN Specifies New Database](#)
- Also see [Application Unable to Connect to Database in Direct Mode](#).

Cannot Connect to the TimesTen Server

You have not correctly identified the system where the TimesTen Server is running.

On a Windows client, select the TimesTen Server in the TimesTen Data Source Setup dialog that is displayed as part of the ODBC Data Source Administrator. To verify the TimesTen Server:

1. On the Windows Desktop, choose **Start > Settings > Control Panel**.
2. Double click the **ODBC** icon. This opens the ODBC Data Source Administrator.
3. Click the **System DSN** tab. This displays the System Data Sources list.
4. Select the TimesTen Client data source. This opens the TimesTen Client DSN Setup dialog.
5. Click **Servers**. This opens the TimesTen Logical Server List.
6. Select the TimesTen Server from the list. This opens the TimesTen Logical Server Name Setup dialog.
7. Verify that the values for the **Network Address** and **Port Number** are correct. If necessary, change the values.

① Note

If you typed the host name or network address directly into the Server Name field of the TimesTen Client DSN Setup, the Client tries to connect to the TimesTen Server using the default port.

If the Network Address and Port Number values are correct, the TimesTen Server may not be running. In the *Oracle TimesTen In-Memory Database Operations Guide*, see Testing Connections for more information about identifying this problem.

On Linux and UNIX, specify the TimesTen Server with the `TTC_Server` connection attribute in the `odbc.ini` file on the client. If the value specified for `TTC_Server` is an actual host name or IP address, the client tries to connect to the TimesTen Server using the default port. In TimesTen, the default port is associated with the TimesTen release number. If the value specified for `TTC_Server` is a logical Server Name, this logical Server Name must be defined in the `ttconnect.ini` file. The `ttconnect.ini` entry for this Server Name must correctly define the host name/IP address and port number on which the TimesTen Server is listening.

If the Network Address and Port Number values are correct, the TimesTen Server may not be running or did not start. See Testing Connections in the *Oracle TimesTen In-Memory Database Operations Guide* for more information about identifying this problem.

TimesTen Server Failed

Check the server's log file. See Creating and Configuring Client DSNs on Linux and UNIX and Managing TimesTen Daemon Attributes in the *Oracle TimesTen In-Memory Database Operations Guide* for information on the TimesTen daemon configuration file.

The maximum number of concurrent IPC connections to the Server of a particular TimesTen instance is 24,999. However, TimesTen has a limit of 2043 connections (direct or client/server) to a single DSN.

Client/server users can change the file descriptor limit to support a large number of connections.

Cannot Find Server DSN

On Linux and UNIX, verify that the Server DSN is defined in the `sys.odbc.ini` file on the system running the TimesTen Server.

On Windows, verify that the Server DSN is defined as a System DSN in the ODBC Data Source Administrator on the system running the TimesTen Server. See Creating and Configuring a Logical Server Name on Windows in the *Oracle TimesTen In-Memory Database Operations Guide*.

TimesTen Server Failed to Load Driver

This error only occurs on Linux or UNIX systems. Open the `sys.odbc.ini` file on the system running the TimesTen Server and locate the Server DSN you are trying to connect. Verify that the dynamic library specified in the `DRIVER` attribute for the Server DSN exists and is executable.

Application Times Out When Accessing TimesTen Server

The default `TimeOut` interval is 60 seconds.

To increase this interval on Linux or UNIX, change the value of the `TTC_Timeout` attribute in the `odbc.ini` file.

TimesTen Client Loses Connection with TimesTen Server

Check to see if the error was due to the Client timing out. Check the TimesTen Server's log to see why the Server may have severed connection with the Client. Use ping to determine if your network is up or try using telnet to connect to the TimesTen Server port number.

Failed to Attach to Shared Memory Segment for IPC

While using shared memory segment (SHM) as IPC, the application may see the following error message from the TimesTen Client ODBC Driver if the application reaches the system-defined per-process file-descriptor-limit.

```
SQLState    = S1000,  
Native Error = 0,  
Message     = [TimesTen][TimesTen 22.1 CLIENT]Failed to attach to shared memory  
segment for IPC. System error: 24
```

This may happen during a connect operation to the Client DSN when the `shmat` system call fails because the application has more open file descriptors than the system-defined per-process file descriptor limit. To correct this problem, you must increase your system-defined per-process file descriptor limit. For more information about file descriptor limits, see System Limits in the *Oracle TimesTen In-Memory Database Reference*.

Thread Stack Overflow When Using Multiple Client Connections

You may receive messages about a segmentation fault that mention a possible thread stack overflow.

If these messages occur, increase the server stack size by one of the following methods:

- Specify the `server_stack_size` option in the `/conf/timesten.conf` file.
- Specify the `ServerStackSize` connection attribute for a specific DSN. This takes precedence over the value in the `/conf/timesten.conf` file.

Increasing the server stack size decreases the number of concurrent connections that can be made before running out of swap space.

See Working with the TimesTen Client and Server in the *Oracle TimesTen In-Memory Database Operations Guide*.

Out of Space When DSN Specifies New Database

You may receive "out of space" messages if you change a DSN to specify a new database while there are existing connections to the original database in a system with multiple client connections.

Close all connections to the original database. This causes a new server process to be created for connections to the database that is now specified in the DSN. Use the `ttStatus` utility to list

the connections for the old database. Alternatively, you can restart the server by using the `ttDaemonAdmin` utility with the `-restartServer` option, which resets all client connections on all DSNs in the instance.

Application Connects or Disconnects Are Slow

This section describes what to check if you encounter slow connects and disconnects to a database. First check your `ramPolicy` setting. If the `ramPolicy` setting is `inUse`, and you are the only user, then at first connect, the entire database is loaded into memory. Consider changing `ramPolicy` to either `always` or `manual`. For more information, see `ttAdmin` in the *Oracle TimesTen In-Memory Database Reference*.

Possible cause	See...
Database is being recovered.	Check If Database Is Being Recovered for information.
ODBC tracing is enabled.	Check ODBC Tracing for information.
There are other possible causes.	API Tracing for information.

Check If Database Is Being Recovered

A slow connect may indicate that a TimesTen database is being recovered. This happens only for a first connect.

Check ODBC Tracing

On Windows platforms, if ODBC tracing is enabled, it can slow connect and disconnect speeds. Double-click **ODBC** in the Control Panel to open the ODBC Data Source Administrator. Select the **Tracing** tab and confirm tracing is disabled. See [Using ODBC Tracing](#) for details.

Application Is Disconnected Unexpectedly

If an application becomes disconnected from a TimesTen database, one of the following events occurs:

- If there was no outstanding transaction, the connection is cleanly removed by the TimesTen daemon. Other existing connections continue processing as if no problem had occurred.
- If there was an outstanding transaction, the transaction is rolled back and the connection is cleanly removed by the TimesTen daemon. Other existing connections continue processing as if no problem had occurred.

This section describes what to check if your application unexpectedly disconnects from the database.

Possible cause	See...
There is an internal application error.	Check for ODBC, JDBC, OCI, Pro*C, and PL/SQL Application Errors for information.
There is failure of a concurrent application thread.	Check for ODBC, JDBC, OCI, Pro*C, and PL/SQL Application Errors for information. Check the User Error Log for information.

Possible cause	See...
If using a client/server connection, the client may have disconnected from the application.	Troubleshooting Client/Server Problems for information.
There is an error in the TimesTen library.	Contact TimesTen Customer Support.

Check for ODBC, JDBC, OCI, Pro*C, and PL/SQL Application Errors

Check for the following types of errors:

- ODBC errors returned by the `SQLError` function
- JDBC errors returned by the `SQLException` class

TimesTen OCI and Pro*C applications, and those that use PL/SQL, report errors using Oracle database error codes instead of TimesTen error codes. The error messages that accompany the error codes may come from the TimesTen error catalog or the Oracle database error catalog.

By default, TimesTen messages and diagnostic information are stored in:

- A user error log that contains error message information. Generally, these messages contain information on actions you may need to take. The default file is `timesten_home/diag/tterrors.log`. For more information on modifying the location of the user error log, see [Error, Warning, and Informational Messages](#) in the *Oracle TimesTen In-Memory Database Operations Guide*.
- A support log containing everything in the user error log plus information of use by TimesTen Customer Support. The default file is `timesten_home/diag/ttmesg.log`. For more information on modifying the location of the user error log, see [Error, Warning, and Informational Messages](#) in the *Oracle TimesTen In-Memory Database Operations Guide*.
- An invalidation file containing diagnostic information when TimesTen invalidates a database. This file provides useful troubleshooting information for TimesTen Customer Support. The invalidation file is created and named based on the value specified by the `DataStore` connection attribute. This connection attribute is not a file name. For example on Linux and UNIX systems, if the `DataStore` connection attribute is `/home/ttuser/AdminData`, the actual invalidation file name has a suffix, `.inval`, `/home/ttuser/AdminData.inval`. For more information on the `DataStore` connection attribute, see [DataStore](#) in the *Oracle TimesTen In-Memory Database Reference*.

It may be helpful to use `ttTraceMon` to generate a level 4 `ERR` trace for the application and review all of the errors messages that are pushed in the TimesTen direct driver. See [ERR Tracing](#) for details.

Check the User Error Log

If a TimesTen application disconnects without returning an ODBC error or any other warning, look through the user error log. See [Using the Logs Generated by the TimesTen Daemon](#) for information.

Application Is Slow

For details on how to maximize the performance of your application and TimesTen database, see:

- TimesTen Database Performance Tuning in the *Oracle TimesTen In-Memory Database Operations Guide*
- ODBC Application Tuning in the *Oracle TimesTen In-Memory Database C Developer's Guide*
- Java Application Tuning in the *Oracle TimesTen In-Memory Database Java Developer's Guide*

This section describes some issues that impair performance.

Possible cause	See...
Using client/server mode.	Consider Connection Mode for information.
Database statistics are outdated.	Update Statistics for Your Tables for information.
Committing transactions too frequently.	Turn Off Autocommit Mode in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
DurableCommits attribute is enabled.	Use Durable Commits Appropriately in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
SQL statements used more than once are not prepared.	Prepare Statements in Advance in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
There is a wrong kind of index, too many indexes, or a wrong size for hash index.	Select the Appropriate Type of Index in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
Use of locks is inefficient.	Size Hash Indexes Appropriately in the Oracle TimesTen In-Memory Database Operations Guide .
Operations are slightly slower when performed on out-of-line columns versus inline.	Verify Lock and Isolation Levels for information.
Materialized view is improperly configured.	Inline and Out-of-Line Columns in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
If replication is used, configuration of replication scheme or network environment may be impacting application.	Performance Implications of Materialized Views and Materialized View Tuning in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .
If TimesTen Cache is used, TimesTen Cache configuration or environment may be impacting application.	Poor Replication or XLA Performance for information.
There are too many table partitions.	Poor Autorefresh Performance for information.
Tracing is unnecessarily enabled for one or more TimesTen components.	Check Partition Counts for Tables for information.
	Check Trace Settings for information.

Consider Connection Mode

Client/server connections are slower than direct connections to TimesTen databases. Driver manager connections can also moderately impact performance. The performance overhead imposed by client/server connections can be significant because of the network latencies involved in all communication with the database.

If your application must run on a different system from the one hosting the database, see Client/Server Tuning in the *Oracle TimesTen In-Memory Database Operations Guide*.

Update Statistics for Your Tables

The TimesTen query optimizer in general is very good at choosing the most efficient query plan. However, it needs additional information about the tables involved in complex queries in order to choose the best plan. By knowing the number of rows and data distributions of column values for a table, the optimizer has a much better chance of choosing an efficient query plan to access that table.

Before preparing queries that will access a TimesTen table, use the `ttOptUpdateStats` procedure to update the statistics for that table. When updating the statistics for a table, you get the best results if you update statistics on your tables after loading them with data, but before preparing your queries. For example, if you update statistics on a table before populating it with data, then your queries are optimized with the assumption that the tables contain no rows (or very few). If you later populate your tables with millions of rows and then execute the queries, the plans that worked well for the situation where your tables contained few rows may now be very slow.

For more information about updating statistics, see *The TimesTen Query Optimizer* in the *Oracle TimesTen In-Memory Database Operations Guide*.

Verify Lock and Isolation Levels

The manner in which multiple applications concurrently access the database can have a major impact on performance.

An application can acquire locks on the entire database, individual tables, and individual rows. Additionally, applications can set an isolation level that determines whether they hold read and update locks until their transactions commit or roll back.

Check the `SYS.SYSTEMSTATS` table, the `SYS.MONITOR` table, or use the `ttXactAdmin` utility to detect whether an application is spending time waiting for locks. See [Check for Deadlocks and Timeouts](#) and [Using the ttXactAdmin Utility](#) for information.

If lock contention is high, you may be able to improve the overall performance of your system by implementing the following:

- Set the `LockLevel` configuration attribute or use the `ttLockLevel` procedure to place locks on rows, rather than on the entire database. Row locking is the default.
- Use the `ttOptSetFlag` procedure to prevent the query optimizer from placing locks on tables. Table locks are sometimes the default, particularly for updates that affect many rows.
- Use read-committed isolation level (`Isolation=1`, the default) for those applications do not require serializable access to the transaction data.

If you see a lot of lock contention, but the above settings are all set to minimize contention, then the contention may be related to the application itself. For example, concurrent threads may be repeatedly accessing the same row. The `ttXactAdmin` utility can sometimes help you detect this sort of contention. Tracing can also be useful in this situation.

For more information about locks and isolation levels, see *Concurrency Control Through Isolation and Locking* in the *Oracle TimesTen In-Memory Database Operations Guide*.

Check Trace Settings

Use `ttTraceMon -e show` as described in [Using the ttTraceMon Utility](#) to confirm tracing is off on all TimesTen components. `ERR` should be set to 1; all other components should be set to 0. Trace levels are preserved when a database is reloaded.

On Windows platforms, confirm that ODBC tracing is disabled. Double-click **ODBC** in the Control Panel to open the ODBC Data Source Administrator. Select the **Tracing** tab and confirm tracing is disabled. See [Using ODBC Tracing](#) for information.

Check Partition Counts for Tables

When a table is created, it has one partition. When you use `ALTER TABLE ... ADD COLUMN` to add new columns, a new partition is added to the table. Adding multiple columns with a single `ALTER TABLE ... ADD COLUMN` statement only adds one partition.

There is a limit of 999 partitions per table. Exceeding this number generates error 8204. An extra read for each new partition slightly degrades performance for each of the new partitions. A high partition count should be avoided. On replicated tables that have multiple partitions, additional space is used for each update on the subscriber side, proportional to the number of partitions. This can result in the subscribers using slightly more perm space than the master.

The partition value for each table is tracked in the `sys16` column of the system table, `SYS.TABLES`. Obtain the partition counts for tables by using the following query:

```
SELECT tblname, sys16 FROM SYS.TABLES;
```

If you discover that a table has too many partitions, do *one* of the following:

- Recreate the table.
- Save and restore the table. Use `ttMigrate -c` to create a migration file. Then restore the table without additional partitions by using `ttMigrate -r -relaxedUpgrade`.

`ALTER TABLE ... DROP COLUMN` does not remove partitions from a table. On replicated systems, all master and subscriber databases must be migrated using the `-relaxedUpgrade` option. Replication does not occur for tables that have different partition structures.

Application Unresponsive, Appears Hung

This section describes what to check if your application is unresponsive and appears to be hung. First use `pstack` or equivalent to determine where your application is spending its time.

Possible cause	See...
There is an internal application error.	Check for ODBC Errors for information.
Connection attributes set in DSN are inconsistent.	Consider Connection Mode for information.
There is excessive lock contention.	Check for Deadlocks and Timeouts for information.

For any possible cause, see [Check Logs and Gather Trace Information](#) for details.

Check Logs and Gather Trace Information

If your application hangs, check the transaction log by using the `ttXactAdmin` utility. See [Using the ttXactAdmin Utility](#) for information.

You can also check the user error log for errors. See [Using the Logs Generated by the TimesTen Daemon](#) for details.

You can also generate a trace log to detect the activities on various TimesTen components. See [Using the ttTraceMon Utility](#) for information.

Check for ODBC Errors

Check the ODBC errors returned by the `SQLError` function in all applications to determine whether one of them has encountered a problem that caused it to hang. Call `SQLError` after each ODBC call to identify error or warning conditions when they first happen. Examples of `SQLError` usage can be found in the demo programs and in Retrieving Errors and Warnings in the *Oracle TimesTen In-Memory Database Error Messages and SNMP Traps*.

If the problem is repeatable, use `ttTraceMon` to generate a SQL trace to determine where the application is hanging. See [SQL Tracing](#) for details. In more extreme cases, it may be helpful to generate a level 4 `ERR` trace for the application and review all of the errors messages that are pushed in the TimesTen direct driver. See [ERR Tracing](#) for details.

Check for Deadlocks and Timeouts

If there is no connect problem, a deadlock or timeout may be the problem. The `SYS.SYSTEMSTATS` table and the `SYS.MONITOR` table records information about deadlocks and timeouts. See [Monitoring the TimesTen System Tables](#) for information on how view the contents of this table. You can also use the `ttXactAdmin` utility to detect the types of locks currently held by uncommitted transactions and the resources on which they are being held.

If a deadlock occurs, the TimesTen subdaemon negotiates the problem by having an application involved in the deadlock generate TimesTen error 6002, "Lock request denied because of deadlock." The error message contains the SQL that the lock holder is running, which can help you diagnose the cause of the deadlock. If your application encounters this error, it should roll back the transaction and then reissue the statements for that transaction. Deadlocks can be caused if your application issues statements in a particular order that results in a circular wait, and can sometimes be prevented by changing the order in which the statements are issued.

An application encounters TimesTen error 6003, "Lock request denied because of timeout," if it cannot acquire a lock within the time period defined by the lock timeout interval set by the `LockWait` attribute in the DSN or by the `ttLockWait` procedure in your application. Upon encountering a timeout error, your application can reissue the statement. Keeping transactions short reduces the possibility of lock timeout errors.

In multithreaded applications, a thread that issues requests on different connection handles to the same database may encounter lock conflict with itself. TimesTen resolves these conflicts with lock timeouts.

Application Unable to Find Previously Created Objects

This section describes what to check if your application cannot locate previously created tables, indexes, sequences or views in the database.

Possible cause	See...
Database is temporary.	Check Temporary DSN Attribute for information.
Overwrite attribute is enabled.	Check Overwrite DSN Attribute for information.
Path name specified in DSN is relative.	Check Path Name to Database for information.

Check Temporary DSN Attribute

Temporary databases (DSN attribute: `Temporary=1`) persist until all connections to the database have been removed. When attempting to access a table in a temporary database and the table does not exist, it is possible that the database in which the table resided in has been dropped.

Check Overwrite DSN Attribute

If the `Overwrite` and `AutoCreate` DSN attributes are enabled and the database already exists, TimesTen drops that database and creates a new one. Any tables that were created in the old database are dropped.

Check Path Name to Database

To ensure that you are always accessing the same database when connecting to a particular DSN, use an absolute database path name instead of a relative one. For example, if the demo database is in the `datastore` directory, specify:

`DataStore=/datastore/demo`

rather than:

`DataStore=demo`

In the latter case, the database path name is relative to the directory where the application was started. If you cannot find a table and you are using a relative database path name, it is possible that the database in which the table resides in does exist but the database (checkpoint and log) files are in a different directory than the one that you are accessing.

See [Specifying Data Source Names to Identify TimesTen Databases](#) in the *Oracle TimesTen In-Memory Database Operations Guide*.

Troubleshooting OCI and Pro*C/C++ Applications

On Windows, if `NLS_LANG` is not set in the environment, the `NLS_LANG` setting is taken from the registry, `HKEY_LOCAL_MACHINE\Software\ORACLE\NLS_LANG`. If `NLS_LANG` is set to an invalid value or if `NLS_LANG` indicates an unsupported character set, either an OCI connection failed error or an `ORA-12705` error is thrown.

See [Globalization Support in TimesTen OCI](#) in the *Oracle TimesTen In-Memory Database C Developer's Guide* for more information on `NLS_LANG`.

Running Out of a Resource

This section describes what to check if TimesTen runs out of resources such as memory space, file system space, file descriptors, and semaphores.

Symptom	See...
Memory consumption seems high.	Operating System Tools and Shared Memory for information.
Running out of memory space.	<ul style="list-style-type: none"> • Operating System Tools and Shared Memory for information. • Check the Amount of Memory Allocated to a Database for information. • Update Query Optimizer Statistics for information. • Check Memory Used by Queries for information. • Out of Memory After Fatal Crash of a Database for information.
Running out of file system space.	Check Transaction Log File Use of File System Space for information.
Running out of transaction log space.	Check If Tracing Is Enabled for information.
Running out of file descriptors.	Check Transaction Log File Use of File System Space for information.
Running out of semaphores.	Increase the Number of Available File Descriptors for information.
Running out of CPU.	Check the Semaphore Limit for information.
	Obtain a stack trace and contact TimesTen Customer Support.

Operating System Tools and Shared Memory

Operating system tools such as `top`, `vmstat`, and `sar` provide statistics about processes and memory usage. The output from these tools can be misleading as an indicator of TimesTen memory consumption because they report shared memory usage for each process but do not report total shared memory usage. Adding together various memory statistics for TimesTen processes overestimates the amount of memory used by TimesTen because shared memory is by definition shared.

Check the Amount of Memory Allocated to a Database

TimesTen uses both permanent and temporary memory regions. The amount of memory allocated for these regions is set by the `PermSize` and `TempSize` attributes in the DSN definition for the database.

When the TimesTen database fills up, it is important to determine whether it is the permanent or the temporary memory region that is filling up. Use the `ttIsql dssize` command to list allocated, in-use, and high water mark sizes for the permanent and temporary memory regions. The `dssize` command selects the following values from `SYS.MONITOR`:

- `PERM_ALLOCATED_SIZE`
- `PERM_IN_USE_SIZE`
- `PERM_IN_USE_HIGH_WATER`
- `TEMP_ALLOCATED_SIZE`
- `TEMP_IN_USE_SIZE`
- `TEMP_IN_USE_HIGH_WATER`

The permanent memory region consists of table and index data, while the temporary memory region consists of internal structures, such as locks, sorting areas, and compiled commands.

Keeping transactions short and making sure there is enough temporary space in the database prevents locks from occupying all of the remaining temporary space. You can also use table locks if transactions are acquiring tens of thousands of row locks.

For tips on how to estimate the size of your database, see [Size Your Database Correctly](#) in the *Oracle TimesTen In-Memory Database Operations Guide*.

Permanent Memory Region Filling Up

Consider whether you can drop any indexes. You may want to look at query plans to see which indexes are actually used. See [Using ttIsql to View and Change Query Optimizer Plans](#) in the *Oracle TimesTen In-Memory Database Operations Guide*. You can also use the `ttRedundantIndexCheck` procedure to discover redundant indexes. The procedure returns suggestions about which indexes to drop.

Use the `ttSize` utility to estimate the amount of memory used by each table in the database. If the amount of data you need to store is too big, you may need to reset the `PermSize` attribute for the database to increase the size of the permanent memory region. Alternatively, you may need to partition your data into several different databases if, for example, you cannot shrink the temporary memory region or create a bigger database because of limits on the memory.

Sometimes when the permanent memory region fills up, copying the data out of the database, deleting all the data, and copying it back in frees up space. This can be done more efficiently by using the `ttMigrate` utility with the `-relaxedUpgrade` option to migrate the data out, destroy and recreate the database, and migrate the data back in.

Finally, you may have to configure the operating system to allow a larger amount of shared memory to be allocated to a process. You may also have to allocate more swap space for virtual memory.

Temporary Memory Region Filling Up

Some commands may be allocating too much space because of out-of-date statistics. See [Update Query Optimizer Statistics](#) for information.

If updating the statistics does not reduce the temporary memory region, disconnect all connections and then reconnect them. Verify that all connections have been disconnected by using the `ttStatus` utility. That frees up all temporary space, but you must reprepare commands.

Diagnose memory usage by queries. See [Check Memory Used by Queries](#) for information.

Update Query Optimizer Statistics

Make sure you have updated the optimizer statistics with the `ttOptUpdateStats` or `ttOptEstimateStats` procedure. To execute some queries, TimesTen must allocate temporary space. The amount of temporary space required is estimated from statistics about the tables used by the query. Without correct statistics, the temporary space required may be underestimated. See [Using the Query Optimizer](#) for information.

Check Memory Used by Queries

You can check the memory that a query uses by observing the high water mark for temporary memory usage. The high water mark represents the largest amount of in-use temporary space used since the high water mark was initialized or reset.

Complete the following tasks:

1. Use the `ttIsql dssize` command to check `TEMP_IN_USE_SIZE` and `TEMP_IN_USE_HIGH_WATER`. Alternatively, you can query the `SYS.SYSTEMSTATS` or the `SYS.MONITOR` table for these values.
2. Call the `ttMonitorHighWaterReset` procedure to reset the `TEMP_IN_USE_HIGH_WATER` to the current value for `TEMP_IN_USE_SIZE`.
3. Execute a query.
4. Use `dssize` to check `TEMP_IN_USE_HIGH_WATER` for peak memory usage for the query.

Out of Memory After Fatal Crash of a Database

Irrecoverable errors, such as errors 846 and 994, invalidate a TimesTen database. However, the database remains in memory, which is only freed after all users have disconnected from the database. If the database is restarted while users are connected to the invalidated database, both old and new instances exist in memory at the same time. In this case, users could receive out-of-memory conditions. To prevent an "Out of memory" error, disconnect all active connections at the time of the irrecoverable error before reconnecting. These irrecoverable errors are rare.

Check Transaction Log File Use of File System Space

TimesTen saves a copy of the database in two checkpoint files, which are stored in the directory specified by the `DataStore` attribute. Each checkpoint file can grow on the file system to be equivalent to the size of the database in shared memory. For each permanent database, you must have enough file system space for the two checkpoint files and for transaction log files.

Transaction log files accumulate in the directory specified by the `LogDir` attribute and are only deleted when checkpoints are performed. If the `LogDir` attribute is not specified in the DSN, transaction log files accumulate in the directory specified by the `DataStore` attribute. The maximum size of your transaction log files is set by the `LogFileSize` attribute.

When a file system fills up with TimesTen data, it is most often due to a build-up of transaction log files. Transaction log files are used for numerous purposes in TimesTen, including transaction rollbacks, backups, and replication. It is important to determine which operation is putting a "hold" on the transaction log files, so that appropriate action can be taken to enable the transaction log files to be purged. This can be done by using the `ttLogHolds` built-in procedure. There are six types of log holds. They are discussed in detail below.

- **Long-running transactions** - TimesTen uses the transaction log to roll back transactions. A log hold is placed for the duration of a transaction. Transactions that are active for a long time result in log file building up if the transaction has written at least one log record. (That is, it is not a read-only transaction.) Commit write transactions with reasonable frequency to avoid significant log file build-up. See Size Transactions Appropriately in the *Oracle TimesTen In-Memory Database Operations Guide* for more information on transaction length.
- **Checkpoint** - If a TimesTen application crashes and the database must be recovered, the checkpoint files and transaction log files are used to recover the data. The "most recent" transaction log files are used -- those written since the checkpoint was done. Transaction log files accumulate during the interval between checkpoints. If checkpoints are done very infrequently, a large number of transaction log files may accumulate, particularly if many changes are made to the database during that interval. See Checkpoint Operations in the *Oracle TimesTen In-Memory Database Operations Guide*.

- **Replication** - TimesTen replication transmits changes from one database to one or more other databases. It does this by reading the transaction log and sending any relevant changes. If replication goes down, the transaction log files build up. See Set the Replication State of Subscribers in *Oracle TimesTen In-Memory Database Replication Guide* for more information on pausing and restarting or resetting replication.
- **Backup** - TimesTen supports an incremental backup facility that uses transaction log files to augment a backup with changes made since the last backup. Transaction log files accumulate during the interval between incremental backups. To avoid a large log build-up, do incremental backups at relatively frequent intervals. If desired, disable incremental backups and do full backups instead.

See Back Up, Restore, and Migrate Data in TimesTen Classic in the *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide* for more information.

- **XLA** - TimesTen's persistent XLA facility reports changes to the database by using transaction log files. Transaction log files are kept until the corresponding transactions have been acknowledged using the `ttXlaAcknowledge` C function. Call `ttXlaAcknowledge` frequently enough to prevent transaction log files building up. See Retrieving Update Records from the Transaction Log in the *Oracle TimesTen In-Memory Database C Developer's Guide*.
- **XA** - TimesTen's XA support uses transaction log files to resolve distributed transactions. If these transactions are not resolved in a timely manner, transaction log files build up. See Distributed Transaction Processing: XA in the *Oracle TimesTen In-Memory Database C Developer's Guide*.

The following attributes are related to file system use.

- The `LogPurge` attribute indicates whether transaction log files that no longer have a hold on them are purged (removed from the file system) or simply archived (renamed). If the `LogPurge` attribute is set to the default value of 0, TimesTen renames transaction log files that it no longer needs by appending the string `.arch` to the name. Once renamed, you must delete the transaction log files manually when they are no longer needed. If transaction log files are not purged, they continue to accumulate space, even when no longer needed by TimesTen.
- The `Preallocate` attribute indicates whether file system space should be reserved for checkpoint files at connect time. This is useful for big databases, to ensure that the file system always has room for the checkpoint files as data is added to the database.

Check If Tracing Is Enabled

When tracing to a file has been enabled, the file may grow so large that a process attempting an operation may exceed the file limits. Tracing always appends to an existing file.

Check the Semaphore Limit

When creating multiple client/server connections to a TimesTen database configured to allow shared memory segment as IPC, you may encounter errors that indicate TimesTen could not create a semaphore. Semaphore limits are platform-dependent.

Duplicate Results from a SELECT Statement

Using read-committed isolation level can lead to duplicates in a result set. A `SELECT` statement selects more or fewer rows than the total number of rows in the table if some rows are added or removed and committed in the range in which the `SELECT` scan is occurring. This may happen when an `UPDATE`, `INSERT` or `DELETE` statement 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.

Index values are ordered. 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 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` statement 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.

For more information about serializable isolation, see *Concurrency Control Through Isolation and Locking* in the *Oracle TimesTen In-Memory Database Operations Guide*.

Cannot Attach PL/SQL Shared Memory

The `PLSQL_MEMORY_ADDRESS` first connection attribute determines the virtual address at which the PL/SQL shared memory segment is loaded into each process that uses the TimesTen direct drivers. Since each operating system platform has different mappings for its address space, the default values for the PL/SQL address space defined in the `PLSQL_MEMORY_ADDRESS` connection attribute are different for each platform, which avoids conflict with operating system mapped address space.

However, if your application overlaps with the PL/SQL mapped address space, you may receive error 8517 "Cannot attach PL/SQL shared memory; `PLSQL_MEMORY_ADDRESS` not valid or already in use." In this case, modify the setting for the `PLSQL_MEMORY_ADDRESS` connection attribute to eliminate the overlap. The reasons for receiving error 8517 can be one of the following:

- User allocated memory already uses that address.
- Some shared memory already uses that address.
- A shared library already uses that address.

If an application accesses two or more TimesTen databases at the same time, you must modify the default setting for the `PLSQL_MEMORY_ADDRESS` attribute in all but one of the TimesTen databases, since the default settings would map the PL/SQL memory address to the same address for all TimesTen databases.

Monitoring and Troubleshooting TimesTen Cache

The following sections in this chapter describe how to monitor and troubleshoot problems you may encounter when using TimesTen Cache:

- [Monitoring Cache Groups with Autorefresh](#)
- [Unable to Create a Cache Group](#)
- [Unable to Start or Stop Cache Agent](#)
- [Unable to Resolve Oracle Service Name](#)
- [Unable to Resolve Connect Identifier](#)
- [Incompatible Oracle Database Server and Client Versions](#)
- [Unable to Validate Oracle Database User Name and Password](#)
- [OCI Initialization Failed](#)
- [Unsupported Data Type Mapping](#)
- [Null Constraint Does Not Match Oracle Database](#)
- [DDL on Cached Oracle Database Tables May Cause Failure](#)
- [Changes Not Visible After Updating Object in Cache Group](#)
- [Load or Refresh Fails](#)
- [Optimize Performance for TimesTen Cache](#)
- [Check Large Batch Jobs on Oracle Database Tables](#)
- [Autorefresh Not Refreshing Cache at Specified Interval](#)
- [Incremental Autorefresh Not Progressing](#)
- [Incremental Autorefresh Becomes Full Autorefresh](#)
- [Poor Autorefresh Performance](#)
- [Declaring NOVALIDATE on Constraints Causes Cache Group Failure](#)
- [AWR Report Shows Lock Contention with DBMS_LOCK](#)

If you are having problems with an AWT cache group, see [Troubleshooting AWT Cache Groups](#) for information.

Note

Error log messages from the cache agent are designated by `CAC` in the message, as shown in several examples in this chapter. These messages are enabled unless you disable the `CACHE` component of the `ttDaemonLog` utility. See `ttDaemonLog` in *Oracle TimesTen In-Memory Database Reference* for information about `ttDaemonLog` options for Windows and Linux or UNIX systems.

Monitoring Cache Groups with Autorefresh

This section includes the following topics:

- [Retrieve Information About Autorefresh Operations](#)
- [Display Information from Change Log Tables](#)
- [Understand Messages About Autorefresh in Support Log](#)
- [Diagnose Autorefresh Failures](#)
- [Diagnose Autorefresh Performance Problems in TimesTen Classic](#)

Retrieve Information About Autorefresh Operations

The `ttCacheAutorefreshStatsGet` built-in procedure returns information about the last ten autorefresh operations on a specified cache group.

The `ttCacheAutorefreshStatsGet` built-in procedure returns information only when the cache agent is running and the autorefresh state is ON or PAUSED. All of the return fields are set to 0 when the cache agent is restarted or the autorefresh state is changed to OFF.

This example uses the `ttCacheAutorefreshStatsGet` built-in procedure for `testcache`, which is a read-only cache group with one table and an incremental autorefresh interval of 10 seconds.

```
Command> call ttCacheAutorefreshStatsGet('user1','testcache');

< 1164260, 2011-07-23 15:43:52.000000, 850280, 44,
0, 75464, 528255, 75464, 310, 110, 6800, 1890912,
12439795, 1890912, 160020, InProgress, 2, 74 >
< 1164260, 2011-07-23 15:43:33.000000, 831700, 43,
13550, 108544, 759808, 108544, 1030, 230, 12290, 1815448,
11911540, 1815448, 160020, Complete, 2, 72 >
< 1164260, 2011-07-23 15:43:12.000000, 810230, 42,
17040, 115712, 809984, 115712, 610, 330, 16090, 1706904,
11151732, 1706904, 146470, Complete, 2, 70 >
< 1164260, 2011-07-23 15:42:52.000000, 790190, 41,
14300, 94208, 659456, 94208, 560, 320, 13410, 1591192,
10341748, 1591192, 129430, Complete, 2, 68 >
< 1164260, 2011-07-23 15:42:32.000000, 770180, 40,
12080, 99328, 695296, 99328, 450, 290, 11340, 1496984,
9682292, 1496984, 115130, Complete, 2, 66 >
< 1164260, 2011-07-23 15:42:12.000000, 750130, 39,
10380, 86016, 598368, 86016, 430, 230, 9720, 1397656,
8986996, 1397656, 103050, Complete, 2, 64 >
< 1164260, 2011-07-23 15:41:52.000000, 730130, 38,
13530, 112640, 700768, 112640, 530, 220, 12780, 1311640,
8388628, 1311640, 92670, Complete, 2, 62 >
< 1164260, 2011-07-23 15:41:32.000000, 710120, 37,
9370, 56320, 326810, 56320, 310, 160, 8900, 1199000,
7687860, 1199000, 79140, Complete, 2, 60 >
< 1164260, 2011-07-23 15:41:22.000000, 700120, 36,
2120, 10240, 50330, 10240, 50, 200, 1870, 1142680,
7361050, 1142680, 69770, Complete, 2, 58 >
< 1164260, 2011-07-23 15:41:12.000000, 690110, 35,
0, 0, 0, 0, 0, 0, 1132440, 7310720, 1132440,
67650, Complete, 2, 56 >
10 rows found.
```

See `ttCacheAutorefreshStatsGet` in the *Oracle TimesTen In-Memory Database Reference* for full details on each column in the output.

In this example, the bolded entries on the first line of the output show that the total number of autorefreshed rows (`totalNumRows = 1890912`) is the same as the total number of autorefreshed root table rows (`totalNumRootTblRows = 1890912`) because there are no child tables.

The number of autorefreshed rows in TimesTen does not necessarily reflect the number of rows updated on the Oracle database. The Oracle database updates may be applied in TimesTen more than once, or multiple Oracle database updates on the same row may be applied as one update in TimesTen.

Display Information from Change Log Tables

TimesTen provides the `cacheInfo` SQL script that gathers information from the change log tables that exist on the Oracle database for cache groups with autorefresh. See *Managing a Cache Environment with Oracle Database Objects* in the *Oracle TimesTen In-Memory Database Cache Guide* for additional information about change log tables.

Run the script as the cache administration user on the Oracle database using SQL*Plus. If you run the script as a different user, it reports that the change log tables do not exist.

The script is in the following location:

```
timesten_home/install/oraclescripts/cacheInfo.sql
```

The `cacheInfo` script displays the following information for each cached table:

```
% cd timesten_home/install/oraclescripts
% sqlplus cacheuser/oracle
SQL> @cacheInfo
*****Autorefresh Objects Information *****
Host name: sys1
Timesten datastore name: /users/OracleCache/alone1
Cache table name: ORATT.ORDERS
Change log table name: tt_06_69245_L
Number of rows in change log table: 100000
Maximum logseq on the change log table: 38
Timesten has autorefreshed updates up to logseq: 38
Number of updates waiting to be autorefreshed: 0
Number of updates that has not been marked with a valid logseq: 0
*****
```

The information returned for each change log table includes the name of the change log table, the name of its corresponding TimesTen cache table, the number of rows in the change log table, and the number of updates in the change log table that have not been automatically refreshed into the cache table.

The log sequence number (logseq) acts as a marker for the autorefresh operation.

Understand Messages About Autorefresh in Support Log

The support log contains messages that show the progress of autorefresh. For example, `testcache` is a read-only cache group with an autorefresh interval of 10 seconds (10,000 milliseconds).

The support log shows when autorefresh starts:

```
15:43:33.96 Info: CAC: 5264: TT47118-5264-5676-refresh03918: Starting autorefresh
number 43 for interval 10000ms
```

The message includes the following information:

- Message number: TT47118
- Timestamp (15:43:33.96)
- Cache agent process ID (5264)
- Thread ID (5676)

The thread ID is important because autorefresh numbers are unique only for a specific interval. Always check both the thread ID and the autorefresh number when you are tracking a specific autorefresh operation.

The support log also contains a longer message that reports information similar to the `ttCacheAutorefreshStatsGet` procedure. 108544 rows were updated in this autorefresh interval, and 1815448 rows have been updated since the cache agent was started. Note that the total number of rows and the total number of root table rows are the same in this message because there is only one table in the cache group. Number refers to the autorefresh number. All times are expressed in milliseconds.

```
15:43:51.81 Info: CAC: 5264: TT47087-5264-5676-refresh04387: Cache agent
refreshed cache group USER1.TESTCACHE: Number - 43, Duration - 13550, NumRows -
108544, NumRootTblRows - 108544, NumOracleBytes - 759808, queryExecDuration -
230, queryFetchDuration - 1030, ttApplyDuration - 12290, totalNumRows - 1815448,
totalNumRootTblRows - 1815448, totalNumOracleBytes - 11911540, totalDuration -
160020
```

Additional messages show that the autorefresh operation completes successfully:

```
15:43:51.81 Info: CAC: 5264: TT47119-5264-5676-refresh04449: Autorefresh
number 43 finished for interval 10000ms successfully
15:43:51.81 Info: CAC: 5264: TT47119-5264-5676-fresher01619: Autorefresh number
43 succeeded for interval 10000 milliseconds
```

Inspect the timestamps to determine whether autorefresh is progressing as expected.

See Managing TimesTen Daemon Attributes in the *Oracle TimesTen In-Memory Database Operations Guide* for information about setting the support log location.

Diagnose Autorefresh Failures

If the `ttCacheAutorefreshStatsGet` built-in procedure shows that the status of an autorefresh operation is **Failed**, check the support log for messages related to the autorefresh operation with number the number shown in the output of the `ttCacheAutorefreshStatsGet` built-in procedure. Look for errors that occurred after the autorefresh operation started.

This row of output from `ttCacheAutorefreshStatsGet` shows a failed autorefresh operation.

```
< 1164260, 2007-08-01 14:56:36.000000, 959350, 9, 0, 0, 0, 0, 0, 0, 1, 7,
1, 50, Failed >
```

The autorefresh number is 9.

The support log shows the start message for autorefresh number 9:

```
14:56:36.10 Info: CAC: 5988: TT47118-5988-4724-refresh03926: Starting
autorefresh number 9 for interval 15000ms
```

The thread ID for autorefresh number 9 is 4724. Look for error messages with this thread ID.

The following messages appear in the support log:

```
14:56:36.10 Info: CAC: 5988: TT47117-5988-4724-refresh03953: Autorefresh thread  
for interval 15000ms is connected to instance inst1 on host host1. Server handle  
231976252  
14:56:36.12 Err : CAC: 5988: TT40018-5988-4724-refresh07567: TimesTen error  
code:5901, msg The Oracle refresh log table, "USER2"."TT_06_81799_L", for base  
table, USER2.READTAB2, cannot be found.  
14:56:36.12 Info: CAC: 5988: TT47055-5988-4724-refresh05559: Autorefresh rolled  
back.  
14:56:36.12 Info: CAC: 5988: TT47119-5988-4724-refresh04458: Autorefresh number  
9 finished for interval 15000ms with error.  
14:56:36.12 Err : CAC: 5988: TT40035-5988-4724-fresher01606: Autorefresh number  
9 failed for cache groups with interval 15000 ms after 10 retries.
```

The error message for thread ID 4724 shows that the change log table, `TT_06_81799_L`, is missing. The introduction to [Autorefresh Not Refreshing Cache at Specified Interval](#) has a table entry that describes what to do in this situation.

Diagnose Autorefresh Performance Problems in TimesTen Classic

On TimesTen Classic, you can use the `ttTraceMon` utility to diagnose autorefresh performance problems. See [AUTOREFRESH Tracing](#) for information.

TimesTen tracing severely impacts application performance and consumes a great deal of file system space if trace output is directed to a file. When you are finished, reset tracing to the default values.

It is also suggested to compare Automated Workload Repository (AWR) Reports when autorefresh is working and when it is not.

Unable to Create a Cache Group

This section describes problems you might encounter when executing the `CREATE CACHE GROUP` statement.

Possible cause	What to do
User does not have the correct Oracle database privileges to create the cache group type.	See Check Oracle Database Privileges .
User has insufficient access to the database.	You must have <code>CACHE_MANAGER</code> privilege to create a cache group.
The internal/external user does not match the Oracle database user.	The TimesTen user name must be the same as the Oracle database user name.

Possible cause	What to do
Cannot connect to the Oracle database.	<p>See:</p> <ul style="list-style-type: none"> • Unable to Resolve Oracle Service Name. • Unable to Resolve Connect Identifier. • Unable to Validate Oracle Database User Name and Password. • Incompatible Oracle Database Server and Client Versions.
Cache administration user ID or password is not set (when trying to create AWT or cache groups with autorefresh).	<p>Also:</p> <ul style="list-style-type: none"> • Check if the Oracle Database needs to be restarted. • Check the network status. <p>See Set Cache Administration User Name and Password.</p>
Data type mapping is unsupported.	See Unsupported Data Type Mapping .
Nullability setting in the Oracle database is different.	See Null Constraint Does Not Match Oracle Database .
Primary key in root table is not specified.	The root table of a cache group must have a primary key. See Defining Cache Groups in the <i>Oracle TimesTen In-Memory Database Cache Guide</i> .

Unable to Start or Stop Cache Agent

This section describes problems you might encounter when starting or stopping the cache agent.

Possible cause	What to do
Cache agent is already running.	See Check Status of Cache Agent .
Cannot locate Oracle Database libraries.	<ul style="list-style-type: none"> • See Check Status of TNS Listener and Oracle Database Server. • Check the permissions on the libraries.
ORACLE_HOME is invalid.	See Check ORACLE_HOME Environment Variable .
Privileges are insufficient.	You must have CACHE_MANAGER privilege to start or stop the cache agent.
OracleNetServiceName is incorrect.	Ensure that the OracleNetServiceName set in your database source name definition matches the Oracle Service Name for the Oracle Database instance that contains the tables to cache in TimesTen.

Check Status of Cache Agent

Check the status of the cache agent:

- On TimesTen Scaleout, use the `ttGridAdmin dbStatus` command. See Monitoring the Status of the Cache Agent Processes in the *Oracle TimesTen In-Memory Database Scaleout User's Guide* for more details.
- On TimesTen Classic, use the `ttStatus` utility. See [Using the ttStatus Utility](#) for more information.

If the cache agent is not running, start it as described in Starting the Cache Agent in the *Oracle TimesTen In-Memory Database Cache Guide*. If attempts to start the cache agent fail, then investigate the possible causes and restart the system before attempting to start the cache agent.

Check ORACLE_HOME Environment Variable

On Linux and UNIX systems, check that the `ORACLE_HOME` environment variable is set correctly for the shell from which you are starting the cache agent and the TimesTen daemon.

See Environment Variables in *Oracle TimesTen In-Memory Database Installation, Migration, and Upgrade Guide* for more information.

Check NLS Environment Variables

NLS environment variables are set in the environment where the TimesTen application is running. Check to see if TimesTen is using the NLS environment variables.

Unable to Resolve Oracle Service Name

If you receive error ORA-12514 indicating "could not resolve service name":

- Use the Oracle Database `TNSPING` utility to verify that the service can be reached.
- Ensure that the `OracleNetServiceName` set in your database source name definition matches the Oracle Service Name for the Oracle Database instance that contains the tables to cache in TimesTen.
- Ensure that there is a service name defined. If it is a Windows Oracle client, use Oracle Net Configuration Assistant to configure a service name. In Oracle Net Configuration Assistant, navigate to Oracle Net Configuration -> Local -> Service Naming, select your Oracle Database server and confirm that there is a service name or a SID that identifies the Oracle Database server. If you add or modify a service name, you may need to restart.

Check the cache administration user name and password on the Oracle database with SQL*Plus to make sure this service name works. For example:

```
%sqlplus cache_admin_user/cache_admin_pwd@OracleHost
```

`cache_admin_user` is the cache administration user name, `cache_admin_pwd` is the cache administration user password, and `OracleHost` is the `OracleNetServiceName` specified in your database source name definition.

Note

Your cache administration user may be different from your regular Oracle database user. See Create the Oracle Database Users and Default Tablespace in the *Oracle TimesTen In-Memory Database Cache Guide*.

- Ensure that there is only one copy of `tnsnames.ora` on your TimesTen system. Also check the permission on `tnsnames.ora`.
- If you run TimesTen on a Linux or UNIX system, check that the `ORACLE_HOME` environment variable is correctly defined. For example:

```
ORACLE_HOME=/products/oracle11g
```

- Check the Oracle Database client and server versions. See [Incompatible Oracle Database Server and Client Versions](#).

Unable to Resolve Connect Identifier

You may receive ORA-12154 "TNS:could not resolve the connect identifier specified" when you try to connect to a database.

This can occur when you are trying to use TimesTen Cache and Oracle Database on the same system and the `TNS_ADMIN` environment variable does not point to the proper `tnsnames.ora` file for Oracle Database. For example, you may have several instances of the Oracle Database running on a laptop.

In a production environment, you typically have TimesTen and Oracle Database running on different systems. In this case, do not reset the `TNS_ADMIN` environment variable to point to a `tnsnames.ora` file on the system where TimesTen is running. The Oracle Database client uses the `TNS_ADMIN` setting to resolve the connection, but the TimesTen main daemon, the cache agent, the Web server, and the replication agent are unaware of the `TNS_ADMIN` setting. TimesTen Cache cannot operate properly when the Oracle Database client and TimesTen use different `tnsnames.ora` files.

On Windows, set the `TNS_ADMIN` environment variable as follows:

1. Open the Start Search, type in `env` and choose "Edit the system environment variables" button.
2. Click the "Environment Variables..." button.
3. Add or edit `TNS_ADMIN` as a system environment variable so that it points to the directory that contains the `tnsnames.ora` file that you want to use. You can include other `tnsnames.ora` files with the `INAME` command inside the `tnsnames.ora` file.

Incompatible Oracle Database Server and Client Versions

If you receive connection timeout errors such as ORA-12170 or ORA-12535, or if you receive ORA-03134 (server version not supported), verify that you are using an Oracle Database client and Oracle Database server whose versions are compatible.

Unable to Validate Oracle Database User Name and Password

This section describes problems you might encounter when using the Oracle database user name and password.

Possible cause	See...
The library environment variable is not set correctly.	Check Library Path Environment Variable for information.
Oracle Database processes are not running.	Check Status of TNS Listener and Oracle Database Server for information.
User does not have the correct Oracle database privileges.	Check Oracle Database Privileges for information.
DSN is incorrectly configured.	Check Database Source Name Definition for information.
There are problems with cache administration user ID or password.	Set Cache Administration User Name and Password for information.

Possible cause	See...
User and system environments are inconsistent.	Check User and System Environment in TimesTen Classic for information.

Check Library Path Environment Variable

Check the library path environment variable on your platform.

Platform	Check Variable
Linux or UNIX	LD_LIBRARY_PATH
Windows	PATH

Check Status of TNS Listener and Oracle Database Server

Try to connect to the Oracle database by using SQL*Plus or use Oracle Enterprise Manager to verify the status.

Check Oracle Database Privileges

From an Oracle SQL*Plus command prompt, list the current Oracle database privileges granted to you by entering:

```
SELECT * FROM SESSION_ROLES;
SELECT * FROM SESSION_PRIVS;
```

Compare the privileges listed against the required privileges for the various TimesTen Cache operations that are specified in Grant Privileges to the Oracle Cache Administration User in the *Oracle TimesTen In-Memory Database Cache Guide*. Contact your Oracle Database Administrator if you require additional privileges.

Check Database Source Name Definition

Take the following actions:

- Confirm you have correctly set the connection attributes.
- On TimesTen Classic, confirm that the DSN definition for TimesTen Cache is a system DSN.
- On TimesTen Classic, confirm that the DSN for TimesTen Cache is defined only once.
- Confirm the Oracle database user name and password. Use SQLPlus and connect to the Oracle database using the same OracleNetServiceName and OraclePWD used in your DSN definition to confirm they are correct.

Reboot TimesTen Machine

If the Oracle Database client was installed and the system has not been restarted, then the TimesTen daemon is still running under the "old" environment before the Oracle Database client install. Restart your system so the TimesTen can start under the "new" environment.

Set Cache Administration User Name and Password

The cache administration user name and password must be set only once in a TimesTen database. However, they must be changed if the TimesTen database is destroyed and recreated or if the cache administration user name is dropped and recreated in the Oracle database.

The cache administration user name and password cannot be changed if the cache agent is running on the TimesTen database or there are cache groups in the database. The cache groups must be dropped before you can change the cache administration user name and password. You must also stop the cache agent before you change the cache administration user name and password, and then restart the cache agent after the user name and password have been changed.

See Registering the Cache Administration User Name and Password in the *Oracle TimesTen In-Memory Database Cache Guide*.

Check User and System Environment in TimesTen Classic

Test to see if the problem is due to differences in user and system environment. This procedure requires two session windows (Command Prompt windows in Windows or shell windows in Linux or UNIX).

1. Stop the TimesTen daemon.
2. In one session window, start the Timesten daemon as a regular user.

On Windows:

```
% timesten_home/install/srv/ttsrv221.exe -d -verbose
```

On Linux or UNIX:

```
% timesten_home/install/srv/timestend -d verbose
```

Some messages will flash by, and then it goes into a wait state.

3. In another session window, try to restart the cache agent.
4. If Step 3 succeeds, then use Ctrl-C on Windows or the `kill` command on Linux or UNIX to stop the TimesTen daemon you started for the other session in Step 2.
5. Compare the user environment and system environment. For example, do both user and system see the same copy of `oci.dll`? Are there any differences in the path name to the `oci.dll` library between the user and system environments?
6. If you detect differences, make the necessary modifications.
7. Restart the system and restart the TimesTen daemon.

Verify Loaded Dynamic Libraries for TimesTen Classic

If you run on a Windows system with Visual C++ installed, verify the loaded dynamic libraries. This works only if you can start the cache agent without autorefresh:

1. Make sure TimesTen is started.
2. Start the cache agent without autorefresh.

```
Command> call ttCacheStart;
Command> create cache group cgl from t1(c1 int not null primary key);
```

3. Open the Windows Task Manager, find process `ttrora221.exe` and highlight it. Right-click on it and select Debug. This brings you into Visual C++ and you should see the loaded DLL in the debug window. See [Unable to Resolve Oracle Service Name](#) for additional information.
4. Load the cache group to force an cache connection from the cache agent:
`Command> load cache group cg1 commit every 100 rows;`
5. Review the loaded DLL in your debug window.

OCI Initialization Failed

Error 5105, "OCI initialization failed," may occur when an operation requires contact with the Oracle database. For example, the error might occur in the following situations:

- Starting the cache agent
- Setting the cache administration user ID or password
- Entering a SQL statement in TimesTen when autocommit=0 and PassThrough=3

Error 5105 contains additional information about its cause:

- OCI cannot find an Oracle database library. See [Check Library Path Environment Variable](#) and check the permissions on the library specified in the error message.
- ORACLE_HOME is invalid. See [Check ORACLE_HOME Environment Variable](#).
- NLS environment variables are set in the environment where the TimesTen application is running. Check to see if TimesTen is using the NLS environment variables.

Unsupported Data Type Mapping

When you try to create a cache group, you may receive the following error:

5115: Unsupported type mapping for column *name*

For example, table *tab* on the Oracle database can be described as follows:

```
COL1      NUMBER(38) NOT NULL
COL2      NUMBER(38)
```

Try to create the cache group as follows:

```
CREATE CACHE GROUP cg FROM tab(col1 CHAR(10) NOT NULL PRIMARY KEY);
```

Error 5119 is displayed and the cache group is not created because the statement attempts to map a column of NUMBER data type to a column of CHAR data type.

See Mappings Between Oracle Database and TimesTen Data Types in the *Oracle TimesTen In-Memory Database Cache Guide*.

Null Constraint Does Not Match Oracle Database

When you try to create a cache group, you may receive the following warning:

Warning 5119: Column *name* has different nullability setting in Oracle

For example, table *tab* on the Oracle database can be described as follows:

```
COL1      NUMBER(38) NOT NULL
COL2      NUMBER(38)
```

Try to create the cache group as follows:

```
CREATE CACHE GROUP cg
  FROM tab(col1 INTEGER NOT NULL PRIMARY KEY, col2 INTEGER NOT NULL);
```

Warning 5119 is displayed because `col2` on the Oracle database does not have a `NULL` constraint, but `col2` in the cache group is defined as `NOT NULL`.

DDL on Cached Oracle Database Tables May Cause Failure

DDL operations that are performed on an Oracle database table that is being cached in TimesTen may cause a failure on the cache group. For example, the user drops a column on the Oracle database table that is being cached in TimesTen. When the cache group is propagated or flushed, TimesTen will update the column that no longer exists in the Oracle database table. When the cache group loads or refreshes, then TimesTen attempts to retrieve data from the column that has been dropped.

The following cache group operations may fail:

- Autorefresh does not occur.
- AWT cache group operations are not propagated or refreshed to or from the Oracle database.
- Cache group load or propagate fails.

If you suspect the cache group operations are not working properly because of a DDL operation on the Oracle database base table, then use DDL tracking to diagnose the issue. DDL tracking saves the change history for all the cached Oracle database tables. The SQL statement and when it was executed are each written to a TimesTen table in the cache administrator user schema on the Oracle database.

For more information on how to create the DDL tracking objects and how to enable DDL tracking for the base table within the Oracle database, see *Tracking DDL Statements Issued on Cached Oracle Database Tables* in the *Oracle TimesTen In-Memory Database Cache Guide*. For details on the built-in procedure used for enabling or disabling DDL tracking, see `ttCacheDDLTrackingConfig` in *Oracle TimesTen In-Memory Database Reference*.

Changes Not Visible After Updating Object in Cache Group

If you modify an object in a cache group and then the changes do not appear on a subsequent SQL statement, then one of the following may have occurred:

- The object was dropped from the Oracle database or was somehow damaged.
- Oracle database was restored or recovered to a time before the object was created.
- The Oracle database was down.
- The user modified the `OracleNetServiceName` DSN or connection attribute after creating the cache group, which points to an Oracle database other than the one that the cache group was created upon.

For example, if the user creates an AWT cache group. Then, the user added rows to a table. When the user performs a `SELECT * FROM` the table, the rows did not appear. The `ttmesg.log` error file does not display an error that the Oracle database is not available. Instead, it displays the following messages:

```
12:09:02.10 Err : REP: 29934: CACHE1:meta.c(904): TT5221: TT5221: Oracle syntax
error in OCIStmtExecute(): ORA-00942: table or view does not exist rc = -1 --
file "bdbStmt.c", lineno 1535, procedure "getOraOutTypesNLengths()"
12:09:02.27 Err : REP: 29934: CACHE1:receiver.c(1978): TT5250: Awt Initialization
Failure. Could not compile meta data sql.
12:09:02.27 Warn: REP: 29934: CACHE1:transmitter.c(6505): TT16060: Failed to read
data from the network. select() timed out
```

To recover, perform the following:

1. Stop all updates to the cache group.
2. If you are using an AWT cache group, then flush the cache group.
3. Recreate the cache group with the drop and create.

Load or Refresh Fails

If the `LOAD CACHE GROUP` or `REFRESH CACHE GROUP` statement fails when you specify `COMMIT EVERY n ROWS` and n is greater than 0, the contents of the target cache group could be in an inconsistent state. Some cache instances may be partially loaded.

Unload the cache group and then load it again. In some situations, it may be easier to drop and recreate the cache group.

Optimize Performance for TimesTen Cache

The following recommendations may optimize performance for the TimesTen Cache:

- Pin the TimesTen Cache meta tables and cache group base tables in the SGA. by executing the `ALTER TABLE table_name CACHE` statement to indicate to the Oracle database that these tables should be stored in the keep portion of the SGA buffer cache. Pinning TimesTen Cache tables in the SGA increases the probability that any given data block needed for a TimesTen Cache refresh operation will be available in the SGA when the refresh is performed and will not force a file system read. This minimizes physical file system reads executed during TimesTen cache refresh operations.
- Pin TimesTen Cache triggers into the shared pool using the `dbms_shared_pool.KEEP` procedure. Pinning triggers into the shared pool for applications where updates to the cache group base tables are infrequent keeps the trigger from having to be reloaded and reparsed. This is not necessary for highly volatile tables where the trigger will be executed frequently and will remain in the shared pool under any circumstances.
- Enable parallel query. For very large base tables with 10 million rows or more, consider using the Oracle database parallel query facility. The primary join query between the log table and the base table is the kind of query which the Oracle database parallel query is designed to handle. When parallel processing is enabled, the parallel query optimizer generates a query plan that enables the original query to be broken into sections to be worked concurrently by different parallel query worker processes. When using parallel query, users should assign a default degree of parallelism of $(2*N)$ to the cache group base table, where "N" is the number of CPUs on the system. Then, experiment to understand what level of parallelism works best for their environment. Experiment with different table structures for base tables, as follows.
 - Use a standard heap table with default degree of parallelism assigned during table creation or by use of the `ALTER TABLE PARALLEL` command. Build an N-partition primary key index against the table.

- Use an N-way partitioned table structure with partition range key based either on the table primary key or, in the case of a concatenated primary key, the high-order column of the primary key. The number of partitions should be set to the degree of parallelism. Use a local primary key index with the same number of partitions.
- Use an N-way hashed partition structure using the primary key as the hash key, a local partitioned primary key index, and both index and table partitions equal to the degree of parallelism. The log table should not be partitioned, as the cardinalities of the log table should never be large enough that a partitioned log table would have any performance benefit. Further, given the continuously increasing value of the log table primary key column, range partitions cannot be used.

 **Note**

Each of these suggestions involve performance tradeoffs, which may not always be beneficial for optimal use. Consider and test each performance suggestion for your own configured environment.

Check Large Batch Jobs on Oracle Database Tables

Customers sometimes run large batch jobs at month-end or year-end on the Oracle database tables that are cached in read-only cache groups with incremental autorefresh. This can cause performance and memory problems for autorefresh operations and replication unless preventative steps are taken.

The tasks described in this section are supported for these components in an TimesTen Cache configuration:

- TimesTen Scaleout
- Active standby pair replication
- Active standby pair with a return service specified
- Active standby pair with a disaster recovery subscriber
- Physical, synchronous Oracle Data Guard
- Oracle RAC

Perform these tasks when large batch jobs must be run on the cached Oracle database tables:

1. Set the autorefresh state to `PAUSED` for cache groups with the `AUTOREFRESH` attribute that are affected by the batch job.
2. Set the autorefresh state to `PAUSED` for cache groups with the `AUTOREFRESH` attribute that are *not* directly affected by the batch job. This ensures that there is a consistent view of the data during batch processing.
3. Run the batch job on the cached Oracle database tables.
4. Make sure all autorefresh change log records have been assigned a valid log sequence number (`logseq`). Call the `cacheInfo.sql` script:

```
% cd timesten_home/install/oraclescripts
% sqlplus cacheuser/oracle
SQL> @cacheInfo
```

Look for the number of updates that have not been marked with a valid log sequence number. Ideally, the number should be zero or small (less than 100) for all tables in cache groups for which the autorefresh state was paused. Consider the following example.

```
*****Autorefresh Objects Information *****
Host name: host1
Timesten datastore name: /scratch/ttuser/ds/mydsn
Cache table name: TTUSER.NOTAFFECTED
Change log table name: tt_05_460491_L
Number of rows in change log table: 1
Maximum logseq on the change log table: 1
Timesten has autorefreshed updates up to logseq: 1
Number of updates waiting to be autorefreshed: 0
Number of updates that has not been marked with a valid logseq: 0
*****
Host name: host2
Timesten datastore name: /scratch/ttuser/ds/mydsn
Cache table name: TTUSER.AFFECTED
Change log table name: tt_05_460489_L
Number of rows in change log table: 100
Maximum logseq on the change log table: 213
Timesten has autorefreshed updates up to logseq: 213
Number of updates waiting to be autorefreshed: 10000
Number of updates that has not been marked with a valid logseq: 0
*****
```

5. For each cache group that was altered in step 1, manually refresh the cache group in parallel mode. Select appropriate values for the transaction size (number of rows committed at a time) and degree of parallelism. For example:

```
REFRESH CACHE GROUP samplecg
  COMMIT EVERY n ROWS PARALLEL m;
  COMMIT;
```

Note that this operation automatically resets the autorefresh state to **ON**.

6. For each cache group that was altered in step 2, set the autorefresh state to **ON**. For example:

```
ALTER CACHE GROUP sampecg2 SET AUTOREFRESH ON;
  COMMIT;
```

Autorefresh Not Refreshing Cache at Specified Interval

The following table shows possible causes for autorefresh problems.

Possible cause	What to do
Cache agent is not started with a cache administration user.	Specify a cache administration user ID and password when starting the cache agent, as shown in Starting the Cache Agent in the <i>Oracle TimesTen In-Memory Database Cache Guide</i> . Also see <code>ttCacheStart</code> in <i>Oracle TimesTen In-Memory Database Reference</i> and Registering the Cache Administration User Name and Password in <i>Oracle TimesTen In-Memory Database Cache Guide</i> .
Object ID of the base table has changed.	See Recover and Reset Autorefresh for Oracle Database Objects for information.
Autorefresh trigger not enabled.	See Recover and Reset Autorefresh for Oracle Database Objects for information.

Possible cause	What to do
Current log sequence number recorded in the TT_version_USER_COUNT table is less than to the maximum log sequence number in the autorefresh log table.	See Recover and Reset Autorefresh for Oracle Database Objects for information.
There is no row in the TT_version_USER_COUNT table with usercount > 0 for every active incremental autorefresh table.	See Recover and Reset Autorefresh for Oracle Database Objects .
Change log table is empty.	See Recover and Reset Autorefresh for Oracle Database Objects .
User count is less than 0 or any TT_version_USER_COUNT log sequence anomalies.	See Recover and Reset Autorefresh for Oracle Database Objects .
Autorefresh log table, trigger, or sequence associated with a cached table does not exist or is not valid.	Check whether the cache agent was started with the correct cache administration user ID. If the cache administration user ID is correct, follow the procedure described in Recover and Reset Autorefresh for Oracle Database Objects .
TT_version_USER_COUNT table is missing.	Check the user error log for messages about "fatal anomalies". This indicates corrupt or missing Oracle database objects.
Current log sequence number in the TT_version_USER_COUNT table changes and is different from the bookmark, and the associated cached table is not refreshed by the next committed autorefresh.	Check whether the cache agent was started with the correct cache administration user ID. If the cache administration user ID is correct, follow the procedure in Recover and Reset Autorefresh for Oracle Database Objects .
If autorefresh occurs at the same time you alter the autorefresh mode, state, or interval, you may potentially cause a lock timeout as both actions modify the same table.	Check the user error log for messages about "fatal anomalies". This indicates corrupt or missing Oracle database objects.
There is a resource problem.	Restart the cache agent. If that does not work, follow the procedure in Recover and Reset Autorefresh for Oracle Database Objects .
	Stop the cache agent before altering the autorefresh mode, state, or interval. See Avoid a Lock Timeout Condition When Modifying Autorefresh .
	Restart the cache agent.

Reset Autorefresh State

Incremental autorefresh does not work if the TRUNCATE statement is used on an Oracle database base table. If TRUNCATE is used on an Oracle database base table, then you must reset autorefresh by using the ALTER CACHE GROUP statement to set the autorefresh state to OFF followed by another ALTER CACHE GROUP to reset the autorefresh state to ON.

Recover and Reset Autorefresh for Oracle Database Objects

If you know or suspect the Oracle database objects used by autorefresh are the cause of the problem, use the following procedure to recreate the Oracle database objects.

1. Use `ALTER CACHE GROUP` to reset the autorefresh state to `OFF` on all cache groups on all databases that have the affected cached table:

```
ALTER CACHE GROUP cache_group_name SET AUTOREFRESH STATE OFF;
```

2. Shut down all cache agents on all affected databases.
3. Check if the user count is zero for each table in the cache group.

On the Oracle database, execute the following statement:

```
SELECT usercount FROM autorefresh_id.tt_version_user_count
  WHERE tablename = 'owner.tablename';
```

If the count is not zero, set the count to zero:

```
UPDATE autorefresh_id.tt_version_user_count SET usercount = 0
  WHERE tablename = 'owner.tablename';
```

4. Start one of the cache agents. The cache agent performs a clean up operation. It displays the following message to the support log after it has completed the cleanup:

```
Cleanup of the Oracle objects completed
```

5. After the cache agent has completed the clean up, use `ALTER CACHE GROUP` to reset the autorefresh state back to `ON`:

```
ALTER CACHE GROUP cache_group_name SET AUTOREFRESH STATE ON;
```

6. Start all other cache agents.

7. Use `ALTER CACHE GROUP` to reset the autorefresh state back to `ON` for all of the affected cache groups on all databases.

Avoid a Lock Timeout Condition When Modifying Autorefresh

A lock timeout may occur if you modify the autorefresh mode, state, or interval with the `ALTER CACHE GROUP` statement at the same time when a scheduled autorefresh starts (scheduled because of the autorefresh interval).

- Autorefresh locks the cache tables and its rows in the `SYS.CACHE_GROUP` table.
- The `ALTER CACHE GROUP SET AUTOREFRESH` statement updates the cache group state in the `SYS.CACHE_GROUP` table.

This results in the following error:

```
Error TT6003: Lock request denied because of time-out.
```

As a workaround, you can:

- Increase the lock timeout with the `ttLockWait` built-in procedure.
- Stop the cache agent before executing the `ALTER CACHE GROUP SET AUTOREFRESH` statement to modify the autorefresh mode, state, or interval.

The following is a TimesTen Classic example that demonstrates setting the autorefresh state to paused.

1. Stop the cache agent.

```
Command> call ttCacheStop;
```

2. Alter the autorefresh state to the desired state. This example sets the autorefresh state to paused.

```
Command> ALTER CACHE GROUP new_customers SET AUTOREFRESH STATE PAUSED;
```

3. Start the cache agent.

Command> call ttCacheStart;

For more information on `ALTER CACHE GROUP`, see `ALTER CACHE GROUP` in the *Oracle TimesTen In-Memory Database SQL Reference*. For more information on the `ttLockWait` built-in procedure, see `ttLockWait` in the *Oracle TimesTen In-Memory Database Reference*.

Incremental Autorefresh Not Progressing

If incremental autorefresh is not progressing, verify that:

- Autorefresh state is ON.
- Cache agent is running.

Inspect the support log for the conditions described in the following table:

Table summary is in the first heading cell.

Condition	What to do
Oracle Database server connection errors or warnings	See Troubleshooting Client/Server Problems for information about resolving connection problems.
Lock timeout errors or warnings on TimesTen	This usually occurs because of an open DDL transaction on the cache group. Commit the DDL transaction so that autorefresh can get the necessary locks.
Insufficient permanent memory region errors on TimesTen	Increase <code>PermSize</code> .
Autorefresh Oracle Database object validations errors or warnings	See Recover and Reset Autorefresh for Oracle Database Objects for information.
Cache agent unexpected exit	Contact TimesTen Customer Support.
Core files in main daemon directory	Contact TimesTen Customer Support.
Warnings about incremental autorefresh becoming full refresh	See Incremental Autorefresh Becomes Full Autorefresh for information.
Warnings that autorefresh has not finished for a long time	The autorefresh transaction can take a long time if many transactions have occurred since the last autorefresh. Note: Cache groups with the same autorefresh interval are autorefreshed in one transaction.

Validate Autorefresh for Oracle Database Objects

The cache agent automatically verifies that the Oracle database objects exist and that they are valid so that autorefresh can progress. In normal operation, you should not see object validation errors or warnings in the user error log. If you see object validation errors, contact TimesTen Customer Support *unless* one of the following conditions has occurred:

- The TimesTen database has been destroyed without using the `DROP CACHE GROUP` statement.
- A customer application inadvertently modifies the objects directly in the Oracle database.
- A DDL operation occurs on the base table on the Oracle database. This disables the trigger that controls autorefresh operations.

The cache group must be recreated if one of the preceding conditions has occurred.

Incremental Autorefresh Becomes Full Autorefresh

On TimesTen Classic, incremental autorefresh can become full autorefresh if the cache administration user tablespace becomes full.

This section includes the following topics:

- [Detect When Incremental Autorefresh Becomes Full](#)
- [Understanding the Cache Administration User Tablespace](#)
- [Diagnose a Full Cache Administration User Tablespace](#)
- [Monitor Usage of a Cache Administration User Tablespace](#)
- [Considerations When a Tablespace Is Full](#)

Detect When Incremental Autorefresh Becomes Full

You can detect when incremental autorefresh becomes full refresh by several methods:

- Check for messages in the support log that indicate full autorefresh operations are occurring. For example:

```
2007-08-08 08:06:51.35 Warn: CAC: 11384: TT47166-11384-1087179104-1Marker01403:  
A full autorefresh will be performed for Incremental autorefresh table  
USER1.READTAB because change log table TT_06_55555_L on Oracle has been  
truncated.
```
- Use the `ttCacheAutorefreshStatsGet` procedure.
 - If autorefresh is *InProgress* for longer than usual, full autorefresh may be occurring.
 - If a much larger number of rows (`autoRefNumRows`) was autorefreshed than usual, full autorefresh may have occurred.

Check the support log for messages about full autorefresh.

Understanding the Cache Administration User Tablespace

TimesTen strongly recommends creating a separate tablespace for the cache administration user. This tablespace is used as the cache administration user's default tablespace. The tablespace contains autorefresh triggers for each Oracle database table, change log tables for each Oracle database table, and other objects that TimesTen needs for each cache administration user. If you do not specify a separate tablespace, then these objects are placed in the Oracle Database system tablespace.

Specify the tablespace when you create the cache administration user on the Oracle database. You can also specify the tablespace after user creation with the `DEFAULT TABLESPACE` clause of the Oracle Database `ALTER USER` statement.

Change log tables for each of the cached Oracle database tables reside in the cache administration user tablespace. For each update on an Oracle database table, one row (a change log record) is inserted into the change log table for that Oracle database table. The size of a change log record in bytes is as follows:

```
size of change log record = size of primary key on Oracle table + 250
```

The number of records in a change log table depends on the update rate on the Oracle database table and on the autorefresh interval on TimesTen. Every 20 seconds, TimesTen

removes change log records that have been applied to all databases that cache the associated Oracle database table.

When change logs are removed, a message similar to the following is displayed in the support log:

```
16:32:26.73 Info: CAC: 5652: TT47112-5652-4756-ogTblGC01036: Garbage collector
deleted 1 rows from TT_06_383270_L where logseq < 1
```

There are options on how to manage what happens when the cache administration user tablespace is filled. See [Considerations When a Tablespace Is Full](#) for more information.

Diagnose a Full Cache Administration User Tablespace

Check for the following conditions if the cache administration user tablespace is full:

- Is the autorefresh state set to PAUSED? Change log records accumulate when the state is PAUSED.
- Has the cache group been created but not loaded? The default autorefresh state for cache group creation is PAUSED.
- Is a cache group being created or is a database being duplicated? Both of these operations temporarily stop clean-up operations on the change log table.
- Are the cache agents on all TimesTen databases running? If a cache agent is not running, change log records accumulate.
- Has a database been abandoned without dropping cache groups with autorefresh in the database? Abandoned databases result from scenarios such as the following:
 - The database is destroyed by `ttDestroy -force`.
 - The application connected to the database with the `Overwrite` connection attribute set to 1, but the cache groups that were in the old database are not recreated.

If the database still exists, connect to the abandoned database and drop the cache group.

Use the `cacheInfo.sql` script to find out how large the change log tables are for each cached Oracle database table. Use the output to verify that the databases are still in use. See [Display Information from Change Log Tables](#).

If the databases are still in use, verify that the cache agents are running.

Compare the autorefresh progress on TimesTen to the maximum log sequence number on the change log table. If TimesTen is behind, then call the `ttCacheAutorefreshStatsGet` procedure to see whether the autorefresh operations are successful. See [Retrieve Information About Autorefresh Operations](#).

If the status is *InProgress* longer than seems reasonable, see [Poor Autorefresh Performance](#).

You may need to decrease the autorefresh interval or increase the size of the cache administration user tablespace.

There are options on how to manage what happens when the cache administration user tablespace is filled. See [Considerations When a Tablespace Is Full](#).

Monitor Usage of a Cache Administration User Tablespace

To monitor the cache administration user tablespace, you can use either Oracle Enterprise Manager alerts or set the TimesTen tablespace threshold parameter.

The cache agent can be configured to periodically monitor the tablespace usage and issue a warning when it exceeds a specified threshold. Set the tablespace threshold percentage with the `TblspaceThreshold` parameter of the `ttCacheConfig` built-in procedure. For example, if you set the `TblspaceThreshold` parameter to 80, then a warning is issued when more than 80% of the tablespace is used.

- If the threshold is set to zero, then no warning is issued. This is the default.
- If the threshold is set between 1 and 99, a warning is issued when the tablespace threshold exceeds that number.
- If the threshold is set to 100, then a warning is issued when the tablespace is full.

For example, to configure for a warning to be issued if the tablespace exceeds 80%, execute the following:

```
call ttCacheConfig('TblspaceThreshold',,'80');
```

For full details of the `ttCacheConfig` built-in procedure, see the `ttCacheConfig` in the *Oracle TimesTen In-Memory Database Reference*.

Considerations When a Tablespace Is Full

With Oracle database tables that are cached in a TimesTen database, you can configure them to use incremental automatic refresh. For these tables, you can specify which one of the following is to occur when the cache administration user's tablespace is full:

- The application performing the DML is to fail. This is the default.
The tablespace full recovery is set to none. The application receives an "Out of Tablespace" error from Oracle Database when the tablespace is full. At that point, the application will need to rollback the transaction.
Setting the tablespace full recovery to none is configured when you set the `Param` parameter to `TblSpaceFullRecovery` and the `Value` parameter to `None` with the `ttCacheConfig` built-in procedure. For example, the following configures `Param` to `TblSpaceFullRecovery` and `Value` to `None` for the `employees` table that is owned by `terry`:

```
call ttCacheConfig('TblSpaceFullRecovery','terry', 'employees', 'None');
```
- Truncate the change log table to free up space and cause a full autorefresh.
When the cache administration user's tablespace is full, any application that is executing DML statements on the autorefresh cached Oracle database tables continues to execute. A trigger executes to free up space for new change log records by deleting existing change log records. This can result in a full automatic refresh on cache groups that have the incremental automatic refresh mode configured. However, if the Oracle database table is not configured for incremental automatic refresh, then no trigger executes.

To set the operation to enable the application to continue and cause an autorefresh, set the `Param` parameter to `TblSpaceFullRecovery` and the `Value` parameter to `Reload` with the `ttCacheConfig` procedure. The user will see stale data until the full autorefresh is complete.

However, even if the user sets the cache configuration parameter `TblSpaceFullRecovery` with the value of `Reload`, the tablespace may not be able to be emptied enough to handle the case of a growing index. Deleting rows from the change log table may not free up enough space for the index that is on the change log table. If the index is growing so fast that it uses all the tablespace to the point where purging the change log tables does not help, then the user's application may receive the following error:

```
ORA-01654: unable to extend index <index> by 128 in tablespace <tblspace>
```

For full details of the `ttCacheConfig` built-in procedure, see the `ttCacheConfig` in the *Oracle TimesTen In-Memory Database Reference*.

Poor Autorefresh Performance

Poor autorefresh performance is usually the result of large autorefresh operations. Use the `ttCacheAutorefreshStatsGet` procedure to check the autorefresh duration and observe whether the status remains *InProgress* for a long time.

Factors that can cause large autorefresh operations include:

- [Incremental Autorefresh Becomes Full Autorefresh](#)
- [Unresponsive TimesTen Database Degrades Autorefresh Performance](#)
- [Refresh of Autorefresh Cache Group Shows Excessive Wait on Resources](#)
- [Abnormally Large Change Log and Base Tables](#)
- [Fragmented Autorefresh Change Log Table Space](#)
- [Degradation in Performance When Autorefresh Interval Is Small](#)
- Large autorefresh interval
- Large number of cache groups with the same interval
- High rate of changes to the Oracle database tables
- The number of generations of child tables in a cache group
- The number of rows in the cached Oracle database tables
- The size of the rows in the cached Oracle database tables

Enable an `AUTOREFRESH` trace to diagnose autorefresh performance problems. See [AUTOREFRESH Tracing](#) for information.

Unresponsive TimesTen Database Degrades Autorefresh Performance

Note

Automatic recovery for TimesTen cache groups only applies to read-only and user managed cache groups that use the `AUTOREFRESH` cache group attribute. In this section, all references to cache groups with autorefresh are read-only and user managed cache groups that use the `AUTOREFRESH` cache group attribute.

If any TimesTen databases containing cache groups with autorefresh are destroyed or no longer in use, TimesTen continues to track autorefresh changes to the Oracle database tables for the TimesTen database for which the cache agent is not running. This causes automatic refresh to cache groups in active TimesTen databases to slow down.

The cache agent is responsible for detecting if a database is unresponsive or no longer in use. You can specify if and how a dead TimesTen database is to be recovered. However, you cannot recover a TimesTen database if all of the Oracle database objects have been removed.

The following sections describe how you can avoid a degraded autorefresh performance for inactive TimesTen databases:

- [Set Cache TimesTen Database Timeout](#)

- [Configure Recovery Method for Cache Groups](#)

Set Cache TimesTen Database Timeout

You can instruct TimesTen to mark the database as dead and no longer accepting updates if the cache agent has not communicated with the Oracle Database server within a specific timeout period.

Set the timeout for the TimesTen database and the recovery method for each cache group with autorefresh by setting the `AgentTimeOut` parameter in the `ttCacheConfig` built-in procedure. The timeout value applies to the all TimesTen databases that use the same cache administration user. You should set the timeout value greater than the time necessary to load the TimesTen database into memory on first connect and start the cache agent. Otherwise, the TimesTen database could be incorrectly marked as dead. For any planned maintenance for the TimesTen instance, you could temporarily set the `AgentTimeOut` value to zero to disable the timeout. For full details of the `ttCacheConfig` built-in procedure, see the `ttCacheConfig` in the *Oracle TimesTen In-Memory Database Reference*.

For example, the following sets the timeout value for the TimesTen database to 6000 seconds or 100 minutes. If the cache agent does not contact the Oracle Database server within a 100-minute period, then the TimesTen database is marked as dead.

```
ttIsql> call ttCacheConfig('AgentTimeOut',,'6000');
```

Configure Recovery Method for Cache Groups

You can recover a TimesTen database and cache groups with autorefresh if they are not synchronizing with the Oracle database. If there is no synchronization, then updates on the Oracle database tables are not automatically refreshed to the corresponding TimesTen cache tables.

You can configure the `DeadDbRecovery` parameter of the `ttCacheConfig` built-in procedure to specify how to recover the synchronization for the TimesTen database and all cache groups with autorefresh. The setting for `DeadDbRecovery` applies to all TimesTen databases that use the same cache administrator user. Set the `DeadDbRecovery` parameter to `Normal`, `Manual` or `None` to describe how TimesTen is to recover the database and all cache groups with autorefresh. The `DeadDbRecovery` setting applies to all TimesTen databases that use the same cache administration user. While TimesTen is recovering the database and its cache groups with autorefresh, there is an autorefresh status for the TimesTen database and the cache groups with autorefresh that describes the recovery status for each of these entities. The TimesTen database can have an automatic refresh status of `Alive`, `Dead` or `Recovering`. The cache groups with autorefresh can have an automatic refresh status of `OK`, `Dead` or `Recovering`. The TimesTen database status changes are linked to changes in the status for the cache groups with autorefresh, as follows:

- If the recovery method is set to `Normal`, then when TimesTen starts a full automatic refresh on an cache group with autorefresh, the cache group's status is set to `Recovering` and the database's status is also set to `Recovering`.
- The TimesTen database's status is only set to `Alive` when all of the cache groups with autorefresh have either been recovered to `OK` or have been dropped.
- When the database status is set to `Dead`, then all of its cache groups with autorefresh are also set to `Dead`.

Note

You can determine the autorefresh status of the TimesTen database and cache groups with autorefresh with the `ttCacheDbCgStatus` built-in procedure, which is described in the `ttCacheDbCgStatus` in the *Oracle TimesTen In-Memory Database Reference*.

When communication between the cache agent and the Oracle Database server is reestablished, TimesTen determines how to recover the cache groups with autorefresh. TimesTen follows the recovery method you configured in the `DeadDbRecovery` parameter in the `ttCacheConfig` built-in procedure. This parameter can be set to one of the following:

- `Normal`: This is the default. The cache groups with autorefresh will each be recovered with a full automatic refresh. After the first full refresh, the cache group is recovered and will incrementally perform autorefresh.

The cache groups with autorefresh within the same automatic refresh interval will be transactionally consistent. Because it is a full refresh, it is not as performant as an incremental refresh.

The autorefresh sets the status to Recovering. When the full autorefresh is completed successfully, the cache group status is set to OK.

- `Manual`: You must manually refresh a cache group with autorefresh to recover it, or unload it if the cache group is dynamic.
- `None`: The cache group with autorefresh will never be recovered by a TimesTen autorefresh. Drop and recreate the cache group to recover it.

The database status changes as the first cache group autorefresh status changes. If there is at least one cache group that is in the process of recovery, then the database status is set to Recovering. Once all cache groups have been recovered, the status of the TimesTen database is marked as Alive.

The following example sets the `DeadDbRecovery` parameter to `Normal` for all cache groups with autorefresh. The dead TimesTen database will be recovered when all of its cache groups with autorefresh have each been recovered with a full automatic refresh.

```
ttIsql> call ttCacheConfig('DeadDbRecovery', , , 'Normal');
```

When TimesTen databases participating in an active standby pair replication scheme contains cache groups, if the autorefresh status of the active master database is Dead and the autorefresh status of the standby master database is Alive, the standby master does not automatically assume the role of the active master. The recovery requires that you manually ensure that the cache and replication agents are executing. The specifics for each situation is as follows:

Table 4-1 Recovery for cache groups involved in active standby replication pair

DeadDbRecovery Setting	Active Master	Standby Master	Resulting Behavior
Normal	Alive	Dead	<p>Make sure that the cache and replication agents are executing on the standby master. Once the cache agent can connect to the Oracle Database, then the status of all cache groups with autorefresh is set to Recovering. This sets the database to Recovering. Once a single cache group has received enough data to resume autorefresh, the status is set to OK. After all cache group are set to OK, the database is set to Alive.</p> <p>Alternatively, you can fail the standby master and rollout a new standby master.</p>
Normal	Dead	Alive	<p>Make sure that the cache and replication agents are executing on the active master. Once the cache agent can connect to the Oracle Database, then the status of all cache groups with autorefresh is set to Recovering. This sets the database to Recovering. Once a single cache group has received enough data to resume autorefresh, the status is set to OK. After all cache group are set to OK, the database is set to Alive.</p> <p>Alternatively, you can fail the active master, switch the standby master as the new active and then rollout a new standby master.</p>
Normal	Dead	Dead	<p>Make sure that the cache and replication agents are executing on both masters. Once the cache agent can connect to the Oracle Database, then the status of all cache groups with autorefresh is set to Recovering. This sets the database to Recovering. Once a single cache group has received enough data to resume autorefresh, the status is set to OK. After all cache group are set to OK, the database is set to Alive.</p> <p>Alternatively, you can rollout new masters.</p>
Manual	Alive	Dead	<p>Make sure that the cache and replication agents are executing on the standby master. Once the cache agent can connect to the Oracle Database, then the status of all cache groups with autorefresh is set to Recovering. This sets the database to Recovering. Once a single cache group has received enough data to resume autorefresh, the status is set to OK. After all cache group are set to OK, the database is set to Alive.</p> <p>Alternatively, you can fail the standby master and rollout a new standby master.</p>
Manual	Dead	Alive	<p>Make sure that the cache and replication agents are executing on the active master. Use a manual refresh to recover the cache groups with autorefresh on the active master. After all cache group are set to OK or have been dropped, the database is set to Alive.</p>
Manual	Dead	Dead	<p>Make sure that the cache and replication agents are executing on the active master. Use a manual refresh to recover the cache groups with autorefresh on the active master. After all cache group are set to OK or have been dropped, the database is set to Alive. Changes are then replicated to the standby master.</p>
None	Alive	Dead	<p>Mark the standby master as failed. Execute <code>ttDestroy</code> utility for the standby master database. Duplicate the active master by executing <code>ttRepAdmin -duplicate</code> utility from the active master.</p>
None	Dead	Alive	<p>Destroy the dead active master with the <code>ttDestroy</code> utility. Recover the dead active master by duplicating the standby master with the <code>ttRepAdmin -duplicate</code> utility.</p>

Table 4-1 (Cont.) Recovery for cache groups involved in active standby replication pair

DeadDbRecovery Setting	Active Master	Standby Master	Resulting Behavior
None	Dead	Dead	Roll out new masters.

Refresh of Autorefresh Cache Group Shows Excessive Wait on Resources

During a refresh of a cache group with autorefresh, there can be excessive buffer busy waits, row lock waits, and deadlocks on updates in the Oracle database, which can negatively affect the throughput performance. When there are multiple deadlocks on updates in the Oracle database involving the autorefresh log tables, the following may appear in the support log:

```
Oracle native error code = 60, msg = ORA-00060: deadlock detected while waiting
for resource
An error occurred while preparing or executing the following Oracle sql
statement: <some statement involving <cache admin user>.TT_##_#####_L where
the # is some number>
```

You can improve your performance by modifying the `INITRANS` and `FREELISTS` settings, which can affect the concurrent inserts into the autorefresh log table and internal maintenance of these tables. The application updating the base table that is being autorefreshed encounters a throughput performance hit when these settings are not appropriately configured.

You can automatically or manually manage these settings as follows:

- Use ASSM tablespace, which automatically manages `FREELISTS`.
- Manually adjust `FREELISTS` and `INITRANS` for the autorefresh log table on the Oracle database.

The following details how to manually modify `INITRANS` and `FREELISTS` for the autorefresh log table on the Oracle database:

1. Retrieve the name of the autorefresh log table that is on the Oracle database.

Under the cache administration user login, execute the SQL*Plus script `cacheInfo.sql` that lists the autorefresh change log table name, along with other items. The following example executes the `cacheInfo.sql` script that lists the autorefresh change log table name as `tt_06_1216726_L`, as shown in bold:

```
SQL> @cacheInfo.sql
*****Autorefresh Objects Information *****
Host name: syst
Timesten datastore name: /users/OracleCache/alone1
Cache table name: ORATT.ORDERS
Change log table name: tt_06_1216726_L
Number of rows in change log table: 1
Maximum logseq on the change log table: 2
Timesten has autorefreshed updates upto logseq: 1
Number of updates waiting to be autorefreshed: 1
Number of updates that has not been marked with a valid logseq: 0
*****
Host name: consyst
Timesten datastore name: /users/OracleCache/alone1
Cache table name: ORATT.ITEMS
Change log table name: tt_06_1279699_L
Number of rows in change log table: 7
Maximum logseq on the change log table: 0
```

```
Timesten has autorefreshed updates upto logseq: 0
Number of updates waiting to be autorefreshed: 5
Number of updates that has not been marked with a valid logseq: 5
*****
```

2. Manually alter the table on the Oracle database. The following example uses the table from the previous example. This example alters the INITRANS and FREELISTS settings for the bar.tt_06_1279699_L table.

```
ALTER TABLE BAR.TT_06_1279699_L INITRANS 10;
ALTER TABLE BAR.TT_06_1279699_L STORAGE(FREELISTS 5);
or
ALTER TABLE BAR.TT_06_1279699_L MOVE STORAGE(FREELISTS 5);
```

3. Alter the INITRANS and FREELISTS settings for the index for this table, which have the same name as the autorefresh change log table with an additional "L" at the end of it. For example, the index for table bar.tt_06_1279699_L is bar.tt_06_1279699_LL.

These settings should be the same as what you set for the autorefresh change log table.

```
ALTER INDEX BAR.TT_06_1279699_LL INITRANS 10;
ALTER INDEX BAR.TT_06_1279699_LL STORAGE(FREELISTS 5);
```

Abnormally Large Change Log and Base Tables

A refresh operation joins the change log table and the base table, which identifies rows needed to be refreshed into TimesTen. The larger the cardinalities of the base table and the change log table, the longer the time necessary to perform this join. Performance degradation may occur if either the change log table or the base table is abnormally large.

The following describes scenarios where the change log table can become abnormally large.

- If the change log table is never purged in configurations where cache groups from multiple DSNs all reference the same base table, it increases in size indefinitely. If one or more of the cache agents for these groups are turned off, those DSNs will not properly refresh their cache groups and the change log tables will not be purged. If the autorefresh state is turned to paused on one of multiple nodes, the other nodes may slow down.
- The change log table can grow abnormally large if some cache agents have been shut down. Resolve this issue by restarting the cache, which will purge all of the backlogged log rows to be purged and all of the cache groups to be synchronized after the completion of the refresh cycle for all cache groups.
- The change log table can be abnormally large if rows inserted into the change log table are never purged and can never be purged by normal processing. This occurs when one or more DSNs are destroyed or rebuilt without first removing the cache groups. The cache group tables on the Oracle database have no information that the cache groups have been destroyed, which corrupts the entire cache group. Rebuild and reinitialize all of the cache groups associated with this base table. Alternatively, never destroy a DSN with cache groups. Instead, always drop the cache groups before destroying a DSN.

Fragmented Autorefresh Change Log Table Space

Change log tables can become fragmented when a high water mark occurs as the result of change logs building up when TimesTen is shut down, for example. If change log tables have become fragmented, you can:

- Coalesce their indexes. This can be done without preventing DML changes to the base tables.

- Perform an online segment shrink. This can be done without preventing DML changes to the base tables.
- Rebuild the change log tables.

Check to see whether space is being wasted:

1. Determine the name of the change log table by running the `cacheInfo.sql` script on the Oracle database.
2. Calculate the size of the change log table. Call the result A. Adapt the name of the change log table in this example.

```
SELECT table_name, ROUND((BLOCKS*8),2) || 'KB' "size"
  FROM user_tables
 WHERE table_name LIKE 'TT_05_%_L';
```

3. Calculate the size of the data in the change log table. Call the result B. Adapt the name of the change log table in this example.

```
SELECT table_name, ROUND((num_rows*avg_row_length/1024),2) || 'KB' "size"
  FROM user_tables
 WHERE table_name LIKE 'TT_05_%_L';
```

4. If $(B/A) * 100$ is greater than 50 percent, then there is at least 40 percent space wasted (assuming a PCTFREE storage parameter set to 10). If there is at least 40 percent space wasted, defragmenting the change log table is recommended.

Perform these steps to defragment the change log table:

1. Alter the cache group to set the autorefresh state to PAUSED.
2. Copy the rows in the change log table to a temporary table.
3. Truncate the change log table.
4. Insert the rows from the temporary table to the change log table.
5. Alter the cache group to set the autorefresh state to ON.

For more information on managing this tablespace, see Managing the Cache Administration User's Tablespace in the *Oracle TimesTen In-Memory Database Cache Guide*.

Degradation in Performance When Autorefresh Interval Is Small

When a relatively short refresh interval, such as a few hundred milliseconds, is combined with a large number of entries in the log table or in the base table, a cache refresh operation does not complete before the next refresh operation is scheduled to begin. In this case, the entries in the log table can be un-marked when the current autorefresh cycle finishes.

Thus, the same rows can be refreshed from the base table to the cache group in the next autorefresh cycle, by which time the rows will be marked. Make sure that the time it for the refresh is greater than the refresh interval. Set the refresh interval to a value where redundant refreshes will not occur.

Declaring NOVALIDATE on Constraints Causes Cache Group Failure

If the Oracle database table on which you want to create the cache group declares NOVALIDATE on columns with primary key, UNIQUE or NOT NULL constraints, the creation of the cache group fails.

Note

This does not apply to any foreign key constraints. However, TimesTen recommends that any matching foreign key is in the enabled VALIDATE state. Your workload performance may be affected when you alter a foreign key column to the enabled VALIDATE state.

TimesTen perceives a NOVALIDATE on a primary key or NOT NULL table column definition as a NULL and, therefore, not qualified as a column on which to build the cache group. Thus, all columns with the primary key, UNIQUE and NOT NULL column constraints must be enabled with the VALIDATE state when creating a cache group from the Oracle database table.

When you create a cache group from an Oracle database table with one or more of these constraints, the following errors are thrown:

5124: Autorefresh/propagate are not allowed on restricted cache group
5168: Restricted cache groups are deprecated
5120: No matching unique index with not null columns, unique key constraint with not null columns, or primary key constraint on table EVENTLOG, cache operations are restricted.

If you receive these errors, you can perform a SELECT statement to verify any existing NOVALIDATE constraints on the Oracle database table. The following SELECT statement shows all constraints on the MyTable table:

```
SQL> select constraint_name, constraint_type, validated, status from
      all_constraints where table_name = 'MyTable';
```

CONSTRAINT_NAME	C	VALIDATED	STATUS
REFID_CONSTRAINT	C	VALIDATED	ENABLED
PKEY_CONSTRAINT	P	NOT VALIDATED	DISABLED

If the table column that is to be the primary key for the cache table is enabled as NOVALIDATE, perform the following steps to enable the column with the VALIDATE state:

1. Enable the NOVALIDATE state for the primary key column.

```
SQL> alter table MyTable modify constraint PKEY_CONSTRAINT
      enable novalidate;
Table altered.
```

```
SQL> select constraint_name, constraint_type, validated, status
      from all_constraints where table_name = 'MyTable';
```

CONSTRAINT_NAME	C	VALIDATED	STATUS
REFID_CONSTRAINT	C	VALIDATED	ENABLED
PKEY_CONSTRAINT	P	NOT VALIDATED	ENABLED

2. Enable the VALIDATE state for the primary key column.

```
SQL> alter table MyTable modify constraint PKEY_CONSTRAINT validate;
Table altered.
```

```
SQL> select constraint_name, constraint_type, validated, status
      from all_constraints where table_name = 'MyTable';
```

CONSTRAINT_NAME	C	VALIDATED	STATUS
-----------------	---	-----------	--------

REFID_CONSTRAINT	C VALIDATED	ENABLED
PKEY_CONSTRAINT	P VALIDATED	ENABLED

AWR Report Shows Lock Contention with DBMS_LOCK

There may be some concern about lock contention when seeing DBMS_LOCK in the Automated Workload Repository (AWR) Report. This wait event is the garbage collector session trying to place a hold on a resource that another garbage collector session from another database has already locked. Thus, only the current garbage collector session waits. The wait for the garbage collector process does not block other processes, except other garbage collectors.

For example, the following shows a contention event in the AWR report:

AWR

Event	Waits	Time (s)	Avg %Total		
			wait	Call	Time Wait Class
eng: UL - contention	2,388	6,997	2930	72.0	Application

In addition, only a small amount of CPU time is used for the garbage collector, as shown in the "SQL ordered by CPU Time" section in the PERF AWR report.

SQL ordered by CPU Time DB/Inst: REM0LNX/REM Snaps: 14976-14977
 -> Resources reported for PL/SQL code includes the resources used by all SQL statements called by the code.
 -> % Total DB Time is the Elapsed Time of the SQL statement divided into the Total Database Time multiplied by 100

CPU Time (s)	Elapsed Time (s)	Executions	CPU per Exec (s)	% DB Time	% Total SQL Id
0	3,508	120	0.00	36.1	0mt5pk2501gph
Module: timestenorad@etcpro01.oracle.com (TNS V1-V3)					
DECLARE v_lockHandle VARCHAR2(200); BEGIN dbms_lock.allocate_unique('ORATT\$ORA_GC1_CACHEADMIN', v_lockHandle); :retval := dbms_lock.request(v_lockHandle, dbms_lock.x_mode, 30, FALSE); END;					
0 3,499 120 0.00 36.0 bb07h2alv817x					
Module: timestenorad@etcpro01.oracle.com (TNS V1-V3)					
DECLARE v_lockHandle VARCHAR2(200); BEGIN dbms_lock.allocate_unique('ORATT\$ORA_DDSMONITOR1_CACHEADMIN', v_lockHandle); :retval := dbms_lock.request(v_lockHandle, dbms_lock.x_mode, 30, FALSE); END;					

Troubleshooting Replication

The following sections in this chapter describe how to troubleshoot problems you may encounter when replicating databases:

- [Unable to Create a Replication Scheme](#)
- [Unable to Alter a Replication Scheme](#)
- [Unable to Start or Stop Replication Agent](#)
- [Replication Does Not Work](#)
- [Replication Unresponsive, Appears Hung](#)
- [Poor Replication or XLA Performance](#)
- [Problems Using ttRepAdmin](#)
- [Problems with Conflict Checking](#)
- [Problems with Parallel Replication](#)

 **Note**

Error log messages from the replication agent daemon are designated by REP in the message, as shown in several examples in this chapter. These messages are enabled unless you disable the REPLICATION component of the ttDaemonLog utility. See ttDaemonLog in *Oracle TimesTen In-Memory Database Reference* for information about ttDaemonLog options.

Unable to Create a Replication Scheme

This section describes what to check if you are unable to use CREATE REPLICATION to create a replication scheme.

Possible cause	What to do
User does not have ADMIN privilege.	You must have ADMIN privilege to use the CREATE REPLICATION or DROP REPLICATION statements.
Database name, host name, or element name is incorrect.	<ul style="list-style-type: none"> • Check the CREATE REPLICATION statement for typographical errors. • See Check Host Names for information. • Use official host names instead of aliases. • The host name should match the value returned by the hostname command on your system and should be used consistently throughout the replication scheme.
The local host is not part of the replication scheme.	Create the replication scheme on a host that will be part of the replication scheme.

Possible cause	What to do
Replication tables defined in the CREATE REPLICATION statement do not exist.	The name, owner, and column definitions of the tables participating in the replication scheme must be identical on both the master and subscriber databases. Use CREATE TABLE to create tables on the database, or use the ttRepAdmin -duplicate utility or the ttRepDuplicateEx C function to duplicate the entire database to be replicated.
There are other problems.	Review the procedures and requirements described in Defining Classic Replication Schemes in the <i>Oracle TimesTen In-Memory Database Replication Guide</i> .

Unable to Alter a Replication Scheme

This section describes what to check if you are unable to use ALTER REPLICATION to alter a replication scheme.

Possible cause	What to do
User does not have ADMIN privilege.	You must have ADMIN privilege to use the ALTER REPLICATION statement.
Replication agent is in Start state.	Most ALTER REPLICATION operations are supported only when the replication agent is stopped (ttAdmin -repStop). Stop the replication agents on both master and subscriber databases, alter the replication scheme on both master and subscriber databases, then restart both replication agents. <ul style="list-style-type: none"> Check ALTER REPLICATION statement for typographical errors. See Check Host Names for information.
Database name, host name, or element name is inconsistent.	Use CREATE TABLE to create a table on the database.
Replication table defined in the ALTER REPLICATION statement does not exist.	Review the procedures and requirements described in Altering a Classic Replication Scheme in the <i>Oracle TimesTen In-Memory Database Replication Guide</i> .
There are other problems.	

Unable to Start or Stop Replication Agent

This section describes what to check if you are unable to start or stop a replication agent.

Possible cause	What to do
User does not have ADMIN privileges.	You must have ADMIN privileges to use the ttAdmin utility or the ttRepStart or ttRepStop procedures to start or stop a replication agent.
TimesTen daemon is not started.	Check the state of the TimesTen daemon. See Check the TimesTen User Error Log for information. If necessary, start the TimesTen daemon as described in Working with the TimesTen Daemon in the <i>Oracle TimesTen In-Memory Database Operations Guide</i> .

Possible cause	What to do
Database does not participate in a replication scheme.	If a database does not participate in a replication scheme, attempts to start a replication agent for that database will fail. Use <code>CREATE REPLICATION</code> to create a replication scheme for the database.

Replication Does Not Work

If you are unable to get replication working between a master and subscriber database, the problem may be one or more of the following:

Possible cause	See...
TimesTen daemon or replication agents are not running.	Check Status of TimesTen Daemon and Replication Agents.
Master and subscriber agents are not communicating.	Check Replication Agents Are Communicating.
Replication is not in Start state.	Check Replication State.
There is an error in the replication scheme.	Check Replication Scheme Configuration.
Owner names for replication scheme and tables are inconsistent.	Check Owner Names.
Replication tables are inconsistent.	Check Consistency between Replicated Tables.

Check Status of TimesTen Daemon and Replication Agents

Use the `ttStatus` utility to confirm the main TimesTen daemon is running and the replication agents are started for all of your master and subscriber databases.

If the TimesTen daemon is running, but the replication agents are not, start the replication agents as described in Starting and Stopping the Replication Agents in the *Oracle TimesTen In-Memory Database Replication Guide*.

If neither the TimesTen daemon or replication agents are running, confirm you have correctly installed TimesTen and the Data Manager service is started.

```
C:\>ttStatus
TimesTen status report as of Mon Aug  6 22:07:53 2012
Daemon pid 5088 port 17000 instance MYINSTANCE
TimesTen server pid 4344 started on port 17002
-----
Data store c:\temp\subscriber1ds
There are 12 connections to the data store
Data store is in shared mode
Shared Memory KEY Global\DBI45b9471c.2.SHM.2 HANDLE 0x280
Type      PID  Context      Connection Name      ConnID
Process   1244  0x00d08fb0  subscriber1ds      1
Replication 4548  0x00aed2f8  LOGFORCE      4
Replication 4548  0x00b03470  TRANSMITTER      5
Replication 4548  0x00b725a8  RECEIVER      6
Replication 4548  0x00b82808  REPHOLD      2
Replication 4548  0x00b98980  REPLISTENER      3
Subdaemon 2752  0x00526768  Worker      2042
Subdaemon 2752  0x0072a758  Flusher      2043
```

```

Subdaemon    2752    0x007308c0  Checkpoint          2044
Subdaemon    2752    0x00736a28  HistGC             2046
Subdaemon    2752    0x067f02f8  Aging              2045
Subdaemon    2752    0x068364a0  Monitor            2047
Replication policy : Manual
Replication agent is running.
Cache agent policy : Manual
-----
Data store c:\temp\masterlds
There are 12 connections to the data store
Data store is in shared mode
Shared Memory KEY Global\DBI45b945d0.0.SHM.6 HANDLE 0x2bc
Type        PID     Context      Connection Name      ConnID
Process     5880   0x00d09008  masterlds          1
Replication 3728   0x00aed570  LOGFORCE           4
Replication 3728   0x00b036e8  TRANSMITTER        5
Replication 3728   0x00b168b8  REPHOLD            3
Replication 3728   0x00b1ca20  REPLISTENER       2
Replication 3728   0x00b22b88  RECEIVER           6
Subdaemon   3220   0x00526768  Worker             2042
Subdaemon   3220   0x0072e768  Flusher            2043
Subdaemon   3220   0x007348d0  Checkpoint         2044
Subdaemon   3220   0x067b0068  Aging              2045
Subdaemon   3220   0x067c0040  Monitor            2047
Subdaemon   3220   0x068404c8  HistGC            2046
Replication policy : Manual
Replication agent is running.
Cache agent policy : Manual
-----
Data store c:\temp\demo
There are no connections to the data store
Replication policy : Manual
Cache agent policy : Manual
-----
End of report

```

This example uses the ttStatus utility to show the output for one master and one subscriber.

```

C:\>ttStatus
TimesTen status report as of Mon Aug  6 22:07:53 2012
Daemon pid 5088 port 17000 instance MYINSTANCE
TimesTen server pid 4344 started on port 17002
-----
Data store c:\temp\subscriberlds
There are 12 connections to the data store
Data store is in shared mode
Shared Memory KEY Global\DBI45b9471c.2.SHM.2 HANDLE 0x280
Type        PID     Context      Connection Name      ConnID
Process     1244   0x00d08fb0  subscriberlds        1
Replication 4548   0x00aed2f8  LOGFORCE           4
Replication 4548   0x00b03470  TRANSMITTER        5
Replication 4548   0x00b725a8  RECEIVER           6
Replication 4548   0x00b82808  REPHOLD            2
Replication 4548   0x00b98980  REPLISTENER       3
Subdaemon   2752   0x00526768  Worker             2042
Subdaemon   2752   0x0072a758  Flusher            2043
Subdaemon   2752   0x007308c0  Checkpoint         2044
Subdaemon   2752   0x00736a28  HistGC            2046
Subdaemon   2752   0x067f02f8  Aging              2045
Subdaemon   2752   0x068364a0  Monitor            2047
Replication policy : Manual
Replication agent is running.

```

```

Cache agent policy : Manual
-----
Data store c:\temp\masterds
There are 12 connections to the data store
Data store is in shared mode
Shared Memory KEY Global\DBI45b945d0.0.SHM.6 HANDLE 0x2bc
Type          PID    Context   Connection Name      ConnID
Process       5880   0x00d09008 masterds           1
Replication   3728   0x00aed570 LOGFORCE          4
Replication   3728   0x00b036e8 TRANSMITTER        5
Replication   3728   0x00b168b8 REPHOLD           3
Replication   3728   0x00b1ca20 REPLISTENER        2
Replication   3728   0x00b22b88 RECEIVER          6
Subdaemon     3220   0x00526768 Worker            2042
Subdaemon     3220   0x0072e768 Flusher           2043
Subdaemon     3220   0x007348d0 Checkpoint        2044
Subdaemon     3220   0x067b0068 Aging             2045
Subdaemon     3220   0x067c0040 Monitor           2047
Subdaemon     3220   0x068404c8 HistGC            2046
Replication policy : Manual
Replication agent is running.
Cache agent policy : Manual
-----
Data store c:\temp\demo
There are no connections to the data store
Replication policy : Manual
Cache agent policy : Manual
-----
End of report

```

This example illustrates that you can run the `ttStatus` utility to check if the replication agent is running. In this example, the replication agent is not running.

```

> ttStatus
TimesTen status report as of Mon Aug  6 22:07:53 2012

Daemon pid 3396 port 15000 instance MYINSTANCE
TimesTen server pid 3436 started on port 15002
-----
Data store c:\temp\subscriberds
There are no connections to the data store
cache agent restart policy: manual
-----
Data store c:\temp\masterds
There are no connections to the data store
cache agent restart policy: manual
-----
End of report

```

This example illustrates that you can run the `ttStatus` utility to check if the daemon and the replication agent are running. In this example, the daemon and the replication agent are not running.

```

> ttStatus
ttStatus: Could not connect to TimesTen daemon: Connection refused

```

Check Replication Agents Are Communicating

Use `ttRepAdmin -receiver -list` to see that the replication agents are communicating with each other. If the `masterds` database is replicating to `subscriberds`, the output should look similar to the following:

```
> ttRepAdmin -receiver -list masterDSN
Peer name      Host name      Port      State  Proto
-----
SUBSCRIBERDS   MYHOST          Auto      Start      10
Last Msg Sent  Last Msg Recv Latency TPS      RecordsPS Logs
-----
0:01:12        -           19.41      5      52      2
```

Check Replication State

You can use the `ttReplicationStatus` built-in procedure to check the state of the subscriber database with respect to its master. If the subscriber is in the `Stop`, `Pause`, or `Failed` state, use the `ttReplicationStatus` procedure to reset the subscriber state to `Start`, as described in Set the Replication State of Subscribers in the *Oracle TimesTen In-Memory Database Replication Guide*.

This examples uses `ttReplicationStatus` to obtain the status of the `subscriberds` database from its master database, `masterDSN`, enter:

```
> ttIsql masterDSN
Command> CALL ttReplicationStatus ('subscriberds');
< SUBSCRIBERDS, MYHOST, 0, pause, 1, 10, REPSCHEME, REPL >
1 row found.
```

To reset state to `Start` call the `ttRepSubscriberStateSet` procedure:

```
Command> CALL ttRepSubscriberStateSet('REPSCHEME', 'REPL', 'SUBSCRIBERDS', 'MYHOST', 0)
Command> CALL ttReplicationStatus ('subscriberds');
< SUBSCRIBERDS, MYHOST, 0, start, 1, 152959, REPSCHEME, REPL >
1 row found.
```

Check Replication Scheme Configuration

This section describes the procedures you can use to confirm the correct configuration of the various components in your replicated system:

- [Check `ttRepAdmin -showconfig`](#)
- [Check TTREP.TTSTORES System Table](#)
- [Check Host Names](#)

Check `ttRepAdmin -showconfig`

Use `ttRepAdmin -showconfig` to confirm the configuration of your replication scheme.

What to look for:

- Are all of the subscriber agents started and reported to be in the `Start` state? If not, reset the agents to the `Start` state. See Set the Replication State of Subscribers in the *Oracle TimesTen In-Memory Database Replication Guide*.
- Do the reported Peer names match the names given in the `DataStore` attributes in the DSN definitions for the replicated databases? Replication does not work if you specified the names given for the `Data Source Name` attributes.
- Is there anything under List of subscribers? If not, confirm the database names you specified in the DSN definition are consistent with those you specified in your replication scheme configuration file.

- Are the Host names correct? If in doubt, see [Check Host Names](#) for information.
- Are the correct table names displayed under Table details? If not, correct the table names in your replication scheme configuration file.

This example confirms the configuration of the replication scheme.

```
> ttRepAdmin -showconfig masterDSN
Self host "MYHOST", port auto, name "MASTERDS", LSN 4/2970276, timeout 120, threshold 0
List of subscribers
-----
Peer name      Host name          Port  State  Proto
-----          -----
SUBSCRIBERDS   MYHOST           Auto   Start   10
Last Msg Sent  Last Msg Recv  Latency  TPS
List of tables and subscriptions
-----
Table details
-----
Table : REPL.TAB
Master Name      Subscriber Name
-----          -----
MASTERDS         SUBSCRIBERDS
```

Check TTREP.TTSTORES System Table

Check the TTREP.TTSTORES table to confirm that replication associates the replication scheme with the local database.

Connect to the database and enter:

```
SELECT * FROM ttrep.ttstores WHERE is_local_store <> 0x0;
< -5193371075573733683, MYHOST, MASTERDS, 01, 0, 0, 4, 0 >
1 row found.
```

There should be exactly one row returned. If more than one row is returned, contact TimesTen Customer Support. If no rows are returned, then none of the hosts returned by the following statement is perceived to be a local system by TimesTen replication:

```
SELECT DISTINCT host_name FROM ttrep.ttstores;
```

It may also be that none of the database names specified in your replication scheme match those specified in your DSN descriptions.

Check Host Names

Some hosts or IP addresses specified in a replication scheme cannot be resolved by the replication agent because:

- Host names or IP addresses specified in the replication scheme are wrong or misspelled.
- Host names or IP addresses cannot be resolved or found by DNS or in the /etc/hosts file
- Entries in the /etc/hosts file are incorrectly ordered in appearance. This error is most common when multiple NICs are used. You must have root privilege to make changes to the /etc/hosts files.

See Configuring the Network in the *Oracle TimesTen In-Memory Database Replication Guide* for details on how to configure DNS and /etc/hosts files for host systems used for replication.

To check if a host name in the replication scheme matches the host name of the local system, write an application to perform these tasks:

1. Use a `gethostname` operating system function call to determine the host name of the running host.
2. Call `gethostbyname` with the output from Step 1.
3. Call `gethostbyname` with the host name specified in the replication scheme.
4. Compare output of Step 2 and Step 3. If there is a match, then the running host is involved in replication. Otherwise, it is not involved in replication.

Check Owner Names

As described in Restrictions and Table Requirements for Classic Replication Schemes and Owner of the Classic Replication Scheme and Replicated Objects in the *Oracle TimesTen In-Memory Database Replication Guide*, the owner names of your replication scheme and your replicated tables must be consistent across all participating databases. This indicates that the replication scheme owner must match in all databases and each table must be owned by the same user in each database.

Check Replication Owner

Check the owner name assigned to your replication scheme by calling the `ttIsql repschemas` command or by listing the contents of the `TTREP.REPLICATIONS` table.

This example shows that the replication scheme name, `REPSCHME`, has a consistent owner name (`REPL`) in the databases on both `SYSTEM1` and `SYSTEM2`.

On `SYSTEM1`:

```
> ttIsql masterDSN
Command> select * from ttrep.replications;
< REPSCHME          , REPL          , C, 0, 0, -1 >
1 row found.
```

On `SYSTEM2`:

```
> ttIsql -connStr "dsn=subscriberDSN"
Command> select * from ttrep.replications;
< REPSCHME          , REPL          , C, 0, 0, -1 >
1 row found.
```

This example shows inconsistent owner names for the replication scheme. This can occur if you omit the owner name from the replication scheme definition and the system uses the Id of the replication scheme creator.

On `SYSTEM1`:

```
> ttIsql masterDSN
Command> select * from ttrep.replications;
< REPSCHME          , SYSTEM1      , C, 0, 0, -1 >
1 row found.
```

On `SYSTEM2`:

```
> ttIsql -connStr "dsn=subscriberDSN"
Command> select * from ttrep.replications;
< REPSCHME          , SYSTEM2      , C, 0, 0, -1 >
1 row found.
```

Check Table Owner

You can check the owner names assigned to the tables in each database by using the `ttIsql` `tables` command.

This example shows that the `TAB` table has a consistent owner name (`REPL`) in the databases on both `SYSTEM1` and `SYSTEM2`.

Output for SYSTEM1	Output for SYSTEM2
<code>REPL.TAB</code>	<code>REPL.TAB</code>

This example shows the `TAB` table with inconsistent owner names, which were automatically assigned for each host.

Output for SYSTEM1	Output for SYSTEM2
<code>SYSTEM1.TAB</code>	<code>SYSTEM2.TAB</code>

Check Consistency between Replicated Tables

By default, TimesTen creates a replication scheme with the `TABLE DEFINITION CHECKING RELAXED` attribute. The `TABLE DEFINITION CHECKING RELAXED` attribute does not require that the column definition of replicated tables be identical. Rather, the `RELAXED` attribute only requires that replicated tables have the same key definition, number of columns, column names, and column data types. If you define the replication scheme with the `TABLE DEFINITION CHECKING EXACT` clause, then the column definitions of replicated tables participating in the replication scheme must be identical on the master and subscriber databases. TimesTen recommends that you use `TABLE DEFINITION CHECKING RELAXED`. For more information, see `Column Definition Options for Replicated Tables` in the *Oracle TimesTen In-Memory Database Replication Guide*.

This output from the user error log shows a mismatch on the number of columns for the subscriber table `TTUSER.MYDSN`.

```
11:37:58.00 Info: REP: 9430: REP1:transmitter.c(4936): TT16136: Sending
definition for table TTUSER.MYDSN (1 column)
11:37:58.00 Info: REP: 9412: REP2:receiver.c(5928): TT16193: Adding definition
for table: TTUSER.MYDSN
11:37:58.00 Info: REP: 9412: REP2:meta.c(5580):TTUSER.MYDSN ptn 0: srcoff 0,
destoff 0, length 8
11:37:58.00 Info: REP: 9412: REP2:meta.c(5580):TTUSER.MYDSN ptn 1: srcoff 8,
destoff 12, length 12
11:37:58.00 Err : REP: 9412: REP2:receiver.c(6203): TT16198: Table definition
mismatch on number of columns for table TTUSER.MYDSN. Local definition: 2;
transmitting peer: 1
11:37:58.00 Err : REP: 9412: REP2:receiver.c(6380): TT16204: Table TTUSER.MYDSN
marked invalid. Will not apply transactions received for it until a valid
definition is received
11:37:58.00 Err : REP: 9412: REP2:receiver.c(7200): TT16078: Table definition
for ID 637068 is invalid (Original failure 11:37:58
REP2:receiver.c(6203): TT16198: Table definition mismatch on number of columns
for table TTUSER.MYDSN. Local definition: 2; transmitting peer: 1)
```

11:37:58.00 Err : REP: 9412: REP2:receiver.c(5002): TT16187: Transaction
1173958671/2; Error: transient 0, permanent 1

Replication Unresponsive, Appears Hung

This table summarizes possible causes when replication is not responding:

Possible cause	See...
There is a failed subscriber.	Check Replication State for information.
Return-receipt timeout period is too long.	Check Return Receipt Timeout Setting for information.

Check Replication State

Use the `ttReplicationStatus` procedure to check state of the subscriber database with respect to its master. If the subscriber is in the Failed state, see Managing Database Failover and Recovery in the *Oracle TimesTen In-Memory Database Replication Guide* for information on how to deal with failed databases.

Use `ttReplicationStatus` to obtain the status of the `subscriberds` database from its master database, `masterDSN`, enter:

```
> ttIsql masterDSN
Command> CALL ttReplicationStatus ('subscriberds');
< SUBSCRIBERDS, MYHOST, 0, failed, 1, 10, REPSCHEME, REPL >
1 row found.
```

Check Return Receipt Timeout Setting

Use the `ttRepSyncGet` procedure to check the return receipt timeout setting. A value of -1 indicates the application is to wait until it receives an acknowledgment from the subscriber. Network latency or other issues might delay receipt of the subscriber acknowledgment. You either address these issues or use the `ttRepSyncGet` procedure to reset the return receipt timeout period.

See Check the Status of Return Service Transactions in the *Oracle TimesTen In-Memory Database Replication Guide* for more information.

Poor Replication or XLA Performance

Most of this section addresses issues that may impact replication performance. Some issues, such as log buffer too small and reading from the transaction log files on the file system, can impact the performance of both replication and XLA applications.

Possible cause	See...
Network is slow.	Check Network Bandwidth for information.
RETURN RECEIPT is used.	Check Use of Return Receipt Blocking for information.
Replication scheme is inconsistent.	Check Replication Configuration for information.
Log buffer is too small.	Check Size of Log Buffer for information.
There are frequent or inefficient file system writes.	Check Durability Settings for information.

Possible cause	See...
Reading from transaction log files on the file system rather than the log buffer.	Check Reads from Transaction Log Files for information.
Rate of conflicts is high.	Conflict Reporting Slows Down Replication for information.

Check Durability Settings

You can improve replication performance by setting `TRANSMIT NODURABLE` on the replication `ELEMENT` to eliminate the flush-log-to-file system operation from the replication cycle. See `Setting Transmit Durability on DATASTORE Element` in the *Oracle TimesTen In-Memory Database Replication Guide* for details.

Enabling the `DURABLE COMMIT` option in your replication scheme also impacts performance. See `DURABLE COMMIT` in the *Oracle TimesTen In-Memory Database Replication Guide* for more information.

Check Network Bandwidth

Replication agents typically communicate over some type of network connection. If replicating over a network slower than 10 MB per second (such as common with a 100 Base-T Ethernet typical in a LAN), you must be careful to match the transaction load to the available bandwidth of the network. see `Network Bandwidth Requirements` in the *Oracle TimesTen In-Memory Database Replication Guide* for details.

Check Use of Return Receipt Blocking

Unless you need receipt confirmation for all your transactions, disable `RETURN RECEIPT BLOCKING`. If you require receipt confirmation for some transactions, then set `RETURN RECEIPT BY REQUEST` and call the `ttRepSyncSet` procedure to enable the return receipt service for specific transactions. See `RETURN RECEIPT BY REQUEST` in the *Oracle TimesTen In-Memory Database Replication Guide* for details.

Note

The performance degradation caused by return-receipt becomes less of an issue when multiple applications (or threads) are updating the database. If you must use return-receipt in a transaction, you can improve the performance of your application by using multiple threads to update the database. Though each thread must block for receipt confirmation, the other threads are free to make updates.

Check Replication Configuration

In addition to return-receipt setting described above, other factors related to the configuration of your replication scheme could impact replication performance. As described in `Making Decisions About Performance and Recovery Tradeoffs` in the *Oracle TimesTen In-Memory Database Replication Guide*, you often have to weigh the ability to efficiently failover and recover a database against replication performance.

For more information about direct replication, see Direct Replication or Propagation in the *Oracle TimesTen In-Memory Database Replication Guide*.

Check Size of Log Buffer

Setting your log buffer too small may impact replication performance. Instead, set the `LogBufMB` DSN attribute to a larger size. For more information on this DSN attribute, see *Setting Connection Attributes for Logging* in the *Oracle TimesTen In-Memory Database Replication Guide*.

Check Reads from Transaction Log Files

In some situations a "log reader," such as a master replication agent 'transmitter' thread or a `ttXlaNextUpdate` call in an XLA application, may not be able to keep up with the update rate of the applications writing to the database. Normally, replication and XLA readers get update records from the log buffer in memory. When the readers fall behind the application update rate, transaction log files can accumulate on the file system until the backlog can be cleared. This forces the readers to read transactions from the transaction log files on the file system, which is much slower. Should you detect reads from the transaction log files, you may want to respond by decreasing the rate of application updates to that sustainable by the log readers.

Applications can monitor whether log readers are obtaining update records from transaction log files on the file system rather than from the log buffer in memory by tracking the `SYS.MONITOR` table entry `LOG_FS_READS`. For example, you can check the value of `LOG_FS_READS` for the database, `MASTERDSN`, with the following `ttIsql` command.

```
% ttIsql -vl -e "select log_fs_reads from monitor; quit;" -connStr dsn=MASTERDSN
```

If the `LOG_FS_READS` counter is increasing, the log readers are falling behind or clearing out a backlog in the transaction log files.

For more complete monitoring of replication progress, create a simple shell script like the following:

```
#!/bin/sh
trap exit 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15
DSN=$1

while [ 1 ] ; do
    date
    ttRepAdmin -receiver -list -connStr dsn=$DSN
    echo -n "Log reads from disk: "
    ttIsql -vl -e "select log_fs_reads from monitor; quit;" -connStr dsn=$DSN
    echo
    ttRepAdmin -bookmark -connStr dsn=$DSN
    sleep 15
done
```

For example, you name the above script `monitorLog` and your replication scheme replicates from the `MASTERDSN` database to the `SUBSCRIBER1DSN` database. You can then check the status of the transaction log by entering:

```
% monitorLog masterdsn
```

This generates output similar to the following:

Mon Aug 2 10:44:40 2004	Peer name	Host name	Port	State	Proto
-----	-----	-----	-----	-----	-----

```

SUBSCRIBER1DSN      MYHOST          Auto   Start   12
Last Msg Sent Last Msg Recv Latency TPS      RecordsPS Logs
-----
00:00:05      -          -1.00      -1      -1      1

Log reads from disk: < 0 >

Replication hold LSN ..... 10/2656136
Last written LSN ..... 10/4015824
Last LSN forced to disk ... 10/3970152

```

The output from the script displays an updated status every 15 seconds until you enter Ctrl-C to exit.

The specific meaning of each value is described in *Display Subscriber Status with the ttRepAdmin Utility* in the *Oracle TimesTen In-Memory Database Replication Guide*. The main interest here is the 'Last Msg Sent' and 'Logs' values. The 'Last Msg Sent' value indicates the elapsed time since the last message was sent by the master to the subscriber and 'Logs' indicates how many transaction log files behind the replication log reader is from the current log insertion point used by the writers (Last written LSN).

Normally the `Logs` value should be 1. A steadily increasing `Logs` value indicates latency is increasing and eventually log reads are satisfied from the file system.

 **Note**

If the `LogBufMB` is larger than the `LogFileSize`, an increase in the 'Logs' value does not necessarily mean the log readers are reading from the transaction log files. This is because the log manager does not allow more than one log file's worth of data to be outstanding before writing it to the file system. After the log manager writes the data, the data remains in the log buffer to be read directly by the log readers. So, when the `LogBufMB` is larger than the `LogFileSize`, the 'Logs' value alone may not be the best measure of whether log readers are reading from memory or from the file system.

The output from the following command displays the number of the transaction log files and the location of the bookmarks set by the log manager.

```
ttRepAdmin -bookmark -connStr dsn=$DSN
```

For more information, see *Monitor Replication with the ttRepAdmin Utility* in the *Oracle TimesTen In-Memory Database Replication Guide*. The difference between the Replication hold LSN and the last written LSN indicates the number of records in the transaction log that have not yet been transmitted to the subscribers. A steady increase in the difference between these values is another indication that replication latency is increasing and log file reads are likely to occur.

In this example, assume the `LogBufMB` is 16MB and the `LogFileSize` is 8MB. The following output indicates the log reader is approximately 1.8 MB behind the capacity of the log buffer and must read from the transaction log files, 14 and 15.

```

Peer name      Host name      Port      State      Proto
-----      -----      -----      -----      -----
SUBSCRIBER1DSN      MYHOST          Auto   Start   12
Last Msg Sent Last Msg Recv Latency TPS      RecordsPS Logs
-----      -----      -----      -----      -----

```

```
00:00:03      -      -1.00      -1      -1      4
Log reads from disk: <20>
Replication hold LSN ..... 14/7007464
Last written LSN ..... 17/465336
Last LSN forced to disk ... 17/456152
```

Problems Using ttRepAdmin

This section includes the following topics:

- [Problems Changing State of Replication Receiver](#)
- [Problems Using ttRepAdmin -duplicate](#)
- [Must Specify Scheme](#)

Problems Changing State of Replication Receiver

If you try to change the state of a replication receiver when replication has been configured to use a secondary IP address, a misconfiguration of the `/etc/hosts` file may cause `ttRepAdmin` to print the following error:

```
Alter replication with 'ALTER REPLICATION...port 0' failed:
TT0907: Unique constraint (REPSTORESIX) violated.
```

This error is caused by replication not recognizing the local database, which can be confirmed by the following query:

```
SELECT * FROM ttrep.ttstores WHERE is_local_store <> 0x0;
```

If this query returns no rows, or returns a row with the main host name of the database set to the result of the `hostname` command rather than the host you specified, you have encountered a configuration problem with the `/etc/hosts` file.

To correct the problem, make sure that the special host name you are using is defined in the `/etc/hosts` file and that there is an IP address in common between your special host name and the result of the `hostname` command.

For example, if the `hostname` command returns `softswitch` and your system has two Ethernet cards with the addresses `10.10.15.136` and `192.168.15.136`, the IP addresses defined for `softswitch` should include both IP addresses.

Problems Using ttRepAdmin -duplicate

Make sure that `Timesten` is installed with the same instance administrator on both the source system and target system. If so, consider these causes:

- [Database Created Before Duplication](#)
- [Wrong Database or Host Name](#)

Database Created Before Duplication

If you connected to your new subscriber DSN before running `ttRepAdmin -duplicate`, the database has already been created. In this situation, `-duplicate` returns:

Error : Restore not done : The datastore already exists.
Unable to restore datastore locally

Confirm the existence of the database by running ttStatus and checking to see if the database is in the returned list. If the new subscriber database exists, destroy it and try ttRepAdmin -duplicate again.

```
> ttDestroy /tmp/newstore
> ttRepAdmin -dsn newstoreDSN -duplicate -name newstore
  -from masterDSN -host "server1"
```

Wrong Database or Host Name

If you have made an error entering the subscriber database name or host name in the replication scheme, you may see something like the following:

Unable to swap datastore locally
No receiver NEWSTORE on SERVER2 found to swap with

If you provide an incorrect host name for a subscriber, you may receive this error:

Problem: ttRepAdmin -duplicate command returns TimesTen error 12080: No subscriber found to swap with.

You also receive information about the subscriber that TimesTen is trying to locate. One common cause is providing an incorrect host name, which must be the exact same name as the host name provided when issuing the CREATE ACTIVE STANDBY PAIR statement. For example, if you created the subscriber with myhost.oracle.com, but only provided myhost in the ttRepAdmin -duplicate, the subscriber will not be found.

Note

If myhost is the local host, use the -localhost argument. You typically need to use the -localhost argument if the local host name does not exactly match the host name provided when creating the replication scheme.

Must Specify Scheme

If you have more than one scheme specified in your TTREP.REPLICATIONS table, some ttRepAdmin commands may return the error:

Must specify -scheme to identify which replication scheme to use

To check the names of the replication schemes used by your database, use the ttIsql utility to connect, and enter:

```
Command> SELECT * from TTREP.REPLICATIONS;
```

This example shows that two replication schemes, REPSCHM1 and REPSCHM2, are assigned to the database associated with subDSN. In this case, it is necessary to use the ttRepAdmin -scheme option.

```
> ttIsql -connStr "dsn=subDSN"
Command> SELECT * from TTREP.REPLICATIONS;
< REPSCHM1      , REPL          , C, 0, 0, -1 >
< REPSCHM2      , REPL          , C, 0, 0, -1 >
2 rows found.
```

```
Command> exit
> ttRepAdmin -dsn subDSN -receiver -list -scheme REPSCHEM1
Peer name          Host name          Port  State  Proto
-----  -----  -----
SUBSCRIBER1        MYHOST           Auto   Start   10

Last Msg Sent Last Msg Recv Latency TPS      RecordsPS Logs
-----  -----  -----  -----  -----
0:01:12          -           19.41      5      52      2
```

Problems with Conflict Checking

This section includes the following topics:

- [Column Invalid for Replication Timestamp](#)
- [Timestamp Does Not Exist](#)
- [Conflict Reporting Slows Down Replication](#)

Column Invalid for Replication Timestamp

When attempting to set `CHECK CONFLICTS` for an element in a `CREATE REPLICATION` statement, you may encounter an error similar to the following.

8004: Column `REPL.TABS.TS` cannot be used for replication timestamp checking if in an index or added by `ALTER TABLE`; and must be `binary(8)` with `NULL` values allowed.

In this situation, check:

- That the timestamp column in the specified table is a nullable column of type `BINARY(8)`. In the above example, the `TS` column in the `REPL.TAB` table should have a type of `BINARY(8)`.
- The timestamp column is defined in the original `CREATE TABLE` statement, rather than added later using `ALTER TABLE`.

Timestamp Does Not Exist

You may receive an error similar to the following:

2208: Column `TS` does not exist in table.

In this situation, confirm that you have specified the correct name for the timestamp column in the `CHECK CONFLICTS` clause and that it exists in the specified table.

Also, make sure the timestamp column is not part of a primary key or index.

Conflict Reporting Slows Down Replication

If there are many conflicts in a short period of time, subscriber performance can slow down because of the reporting requirements. You can use store attributes in the `CREATE REPLICATION` or `ALTER REPLICATION` statements to suspend and resume conflict reporting at specified rates of conflict:

- `CONFLICT REPORTING SUSPEND AT rate`
- `CONFLICT REPORTING RESUME AT rate`

Information about conflicts that occur while reporting is suspended cannot be retrieved.

See Reporting Conflicts in the *Oracle TimesTen In-Memory Database Replication Guide*.

Problems with Parallel Replication

If you have configured parallel replication, then you have additional TimesTen replication threads running on your system. Be sure that you have enough CPU cores to accommodate the increased number of replication threads.

If you have configured user-defined track based parallel replication (`ReplicationApplyOrdering = 1` and `ReplicationParallelism >= 2`), and you are seeing permanent errors on the receiver side that include constraint violations or row lookup failures, then possible causes include:

- Application defined dependencies are incorrect
- Operations for the same keys or keys of tables in a constraint relationship are written to different tracks.

Check the value of the general connection attribute, `ReplicationTrack`. All transactions executed by the connection are assigned to the track as defined in the `ReplicationTrack` attribute setting. Make sure the workload for each track is suitable. Use the `ALTER SESSION` statement to change the replication track. For more information, see `ALTER SESSION` in the *Oracle TimesTen In-Memory Database SQL Reference*. For more information on configuring parallel replication, see Configuring Parallel Replication in the *Oracle TimesTen In-Memory Database Replication Guide*.

Troubleshooting AWT Cache Groups

Creating an asynchronous writethrough (AWT) cache group automatically creates a replication scheme that enables the database to communicate with the Oracle database. You must start the replication agent after you create an AWT cache group and before you start the cache agent.

This chapter contains the following troubleshooting sections:

- [Replication Agent Not Started](#)
- [AWT Performance Monitoring](#)
- [Possible Causes of Poor AWT Performance](#)
- [Oracle Database Errors Reported by TimesTen for AWT](#)

Replication Agent Not Started

The replication agent must be started in order for the AWT cache group to work. If the replication agent is not starting, use the information in [Troubleshooting Replication](#) to troubleshoot the replication agent.

If you are unable to get replication working, the problem may be solved by one or more of the following sections:

Possible Cause	See...
Replication agent will not start.	Unable to Start or Stop Replication Agent for information.
TimesTen daemon or replication agents are not running.	Check Status of TimesTen Daemon and Replication Agents for information.
Replication agents are not communicating.	Check Replication Agents Are Communicating for information.
Replication is not in Start state.	Check Replication State for information.

AWT Performance Monitoring

You can monitor the performance of asynchronous writethrough (AWT) cache groups to determine how much time is spent performing tasks in the AWT workflow.

The following sections describe how to enable the monitoring and the tools for viewing the results.

- [Enable AWT Performance Monitoring](#)
- [Display AWT Performance Results with the ttRepAdmin Utility](#)
- [Use System Tables to Monitor AWT Operations](#)

Enable AWT Performance Monitoring

Use the `ttCacheAWTMonitorConfig` built-in procedure to enable monitoring for AWT cache groups. Ensure that the replication agent is running before calling `ttCacheAWTMonitorConfig`.

The following example enables monitoring and sets the sampling frequency to 16. A sampling factor of 16 is recommended for accuracy and performance.

```
Command> CALL ttCacheAWTMonitorConfig('ON',16);
< ON, 16 >
1 row found.
```

Display AWT Performance Results with the `ttRepAdmin` Utility

Use the `ttRepAdmin` utility with the `-awtmoninfo` and `-showstatus` commands to display the monitoring results. The AWT monitoring statistics include the following:

- TimesTen processing time: The total number of milliseconds spent in processing AWT transaction data since monitoring was enabled.
- Oracle Database bookmark management time: The total number of milliseconds spent in managing AWT metadata on Oracle Database since monitoring was enabled.
- Oracle Database execute time: The total number of milliseconds spent in OCI preparation, binding and execution for AWT SQL operations since monitoring was enabled. This statistic includes network latency between TimesTen and Oracle Database.
- Oracle Database commit time: The total number of milliseconds spent in committing AWT updates on Oracle Database since monitoring was enabled. This statistic includes network latency between TimesTen and Oracle Database.
- Time since monitoring was started.
- Total number of TimesTen row operations: The total number of rows updated in AWT cache groups since monitoring was enabled.
- Total number of TimesTen transactions: The total number of transactions in AWT cache groups since monitoring was enabled.
- Total number of flushes to the Oracle database: The total number of times that TimesTen data has been sent to the Oracle database.
- The percentage of time spent on TimesTen processing, Oracle Database bookmark management, Oracle Database execution and Oracle Database commits.

For example:

```
ttRepAdmin -showstatus -awtmoninfo myDSN

[other -showstatus output]
...
AWT Monitoring statistics
-----
TimesTen processing time : 8443.424000 millisecs (44.064958 )
Oracle execute (SQL execution) time : 8930.320000 millisecs (46.605994 )
Oracle execute (PL/SQL execution) time : 0.000000 millisecs (0 )
Oracle commit time : 1787.568000 millisecs (9.329048 )
Time since monitoring was started: 21954.410000 millisecs
CacheAwtMethod mode : 0
Cache-connect Operational Stats :
    SQL Operations sent to Oracle : 143556
    Number of update operations : 58
```

```

Number of update batches      : 58
Number of insert operations  : 143498
Number of insert batches     : 1146
Number of delete operations  : 0
Number of delete batches     : 0
Total number of batches sent: 1204
Number of bytes sent        : 4769094
Number of TimesTen Transactions sent to Oracle (includes retries) :
143556
Number of retries on TimesTen due to errors on Oracle : 0
Number of round trips to Oracle (includes executes, commits and
rollbacks) : 2290
Number of commits on Oracle : 1086
Number of rollbacks on Oracle : 0
Number of rxbatchs: 244
Number of rxskips: 0

Peer name          Host name          Port  State Proto
-----  -----
_ORACLE           MYHOST           Auto  Start   27

Last Msg Sent  Last Msg Recv  Latency  TPS      RecordsPS Logs
-----  -----
00:00:00        -          2.35    6311      6311      1

```

Use System Tables to Monitor AWT Operations

To determine whether asynchronous writethrough operations are keeping up with the rate of updates on cache tables in AWT cache groups, query the `LAST_LOG_FILE`, `REPHOLD_LOG_FILE` and `REPHOLD_LOG_OFF` columns of the `SYS.MONITOR` system table. If the difference between the value in the `LAST_LOG_FILE` column and the value in the `REPHOLD_LOG_FILE` column is increasing over time, and the value in the `REPHOLD_LOG_OFF` column is increasing slowly or not changing, then the tables are being updated at a faster rate than the updates are being replicated.

Then run a `ttRepAdmin -receiver -list` utility command and find the row where `_ORACLE` is in the `Peer name` field. View the values in the `Last Msg Sent`, `Last Msg Recv`, `Latency` and `TPS` fields within the same row to determine if the replication activity that is falling behind is asynchronous writethrough operations.

For more information about monitoring log files, see Monitoring Accumulation of Transaction Log Files in *Oracle TimesTen In-Memory Database Operations Guide*.

The `SYS.SYSTEMSTATS` system table has a variety of statistics about AWT performance. See `SYS.SYSTEMSTATS` in *Oracle TimesTen In-Memory Database System Tables and Views Reference*.

Possible Causes of Poor AWT Performance

This section addresses issues that may degrade AWT cache group performance.

Possible cause	See...
Network is slow.	Check Network Bandwidth for information.
Log buffer is too small.	Check Size of Log Buffer for information.
There are frequent or inefficient file system writes.	Check Durability Settings for information.

Possible cause	See...
Reading from transaction log files on the file system instead of the log buffer.	Check Reads from Transaction Log Files for information.
There are problems on the Oracle database.	Problems on the Oracle Database for information.
Replication agent is down.	Unable to Start or Stop Replication Agent or Replication Does Not Work for information.
UNLOAD CACHE GROUP cannot complete if row updates are not propagated to the Oracle database.	UNLOAD CACHE GROUP Requires Row Updates to Complete Propagation for information.

Problems on the Oracle Database

If an AWT cache group cannot propagate updates to the Oracle database, the causes can include network failure, an Oracle database failure, I/O bottlenecks and CPU bottlenecks.

AWT propagation to the Oracle database can also fail because the replication agent is stopped. See [Unable to Start or Stop Replication Agent](#) and [Replication Does Not Work](#) for information.

UNLOAD CACHE GROUP Requires Row Updates to Complete Propagation

Execution of the UNLOAD CACHE GROUP statement for an AWT cache group waits until updates on the rows have been propagated to the Oracle database. However, the UNLOAD CACHE GROUP statement may fail if data propagation to the Oracle database is stopped for any reason, such as network failure, the Oracle database is down, or the replication agent is stopped.

Oracle Database Errors Reported by TimesTen for AWT

The following sections describe both permanent and transient Oracle Database errors that can be reported by TimesTen for an AWT cache group:

- [Permanent Oracle Database Errors Reported by TimesTen](#)
- [Transient Oracle Database Errors Reported by TimesTen](#)

Permanent Oracle Database Errors Reported by TimesTen

Insert, update, or delete errors that occur while applying changes to Oracle Database are saved in an error file located in the database directory with the following name:

DataSourceName.awterr

Errors reported to this file are *permanent* errors. TimesTen does not retry the transaction. The errors may be reported in the AWT error file long after the commit to TimesTen occurs.

The format of the messages in the AWT error file is similar to those generated for conflict and transaction errors in replication. Oracle Database error messages are also reported in the support log and the user log. The Oracle Database errors are prefixed with ORA-. The TimesTen error messages are prefixed with TT..

If a constraint violation occurs when a cache group update is propagated to the Oracle database, the message in the AWT error file is similar to the following:

```
Error occurred 14:48:55 on 03-22-2011
Datastore: c:\temp\cgDSN
Oracle Id: system1
Transmitting name: cgDSN
Error message:
    TT5210: Oracle unique constraint violation error in OCIStmtExecute():
    ORA-00001: unique constraint (GUSER.SYS_C00357240) violated rc = -1 -- file
    "bdbTblH.c", lineno 1205, procedure "ttBDBstmtForce()"
    TT5025: Commit failure in Oracle. Transaction must be rolled back in TimesTen.
-- file "bdbConnect.c", lineno 885, procedure "ttBDBxact()"

Operation that caused the error:
Insert into table TESTUSER.T1 <9,1000>

Failed transaction:
Insert into table TESTUSER.T1 <9, 1000>
End of failed transaction
```

If an object that TimesTen has placed on the Oracle database is dropped, the message in the AWT error file is similar to the following:

```
May 04 18:12:36 HOST1 TimesTen Replication 11.2[2136]:
[Err ] DEFAULT:meta.c(639):
TT16062: Failed to compile command:
select p.commit_timestamp, p.commit_seqnum, p.protocol from owner1.TT_06_REPPEERS
p where p.replication_name = :rname and p.replication_owner = :rowner and
p.tt_store_id = :oid and p.subscriber_id = :sid

May 04 18:12:36 HOST1 TimesTen Replication 11.2[2136]:
[Err ] DEFAULT:meta.c(639):
TT5221: TT5221: Oracle syntax error in OCIStmtExecute():
ORA-00942: table or view does not exist rc = -1 -- file "bdbStmt.c", lineno 1041,
procedure "getOraOutTypesNLengths()"
```

In this example, the `TT_06_REPPEERS` table does not exist. To recover from this error, perform the following tasks:

1. Stop the replication agent.
2. Drop and recreate the cache group.
3. Restart the replication agent.

Transient Oracle Database Errors Reported by TimesTen

The support log for databases with AWT cache groups may contain Oracle Database errors if the replication agent encounters a problem on the Oracle database. If the replication agent encounters one of these errors, AWT rolls back the transaction and retries it. If the support log becomes full, the oldest messages are deleted and replaced by new messages.

The Oracle Database errors in the support log are considered *transient* because AWT retries the transaction.

Some transient errors indicate an underlying problem on the Oracle database must be solved before AWT operations can continue. For example:

```
ORA-01536: space quota exceeded for tablespace
ORA-01034: ORACLE not available
```

After the underlying problem has been fixed, AWT retries the operation.

The following Oracle Database errors are transient:

ORA-00018: maximum number of sessions exceeded
ORA-00019: maximum number of session licenses exceeded
ORA-00020: maximum number of processes (%s) exceeded
ORA-00025: failed to allocate %s
ORA-00028: your session has been killed
ORA-00038: Cannot create session: server group belongs to another user
ORA-00051: timeout occurred while waiting for a resource
ORA-00052: maximum number of enqueue resources (%s) exceeded
ORA-00053: maximum number of enqueues exceeded
ORA-00054: resource busy and acquire with NOWAIT specified
ORA-00055: maximum number of DML locks exceeded
ORA-00057: maximum number of temporary table locks exceeded
ORA-00058: DB_BLOCK_SIZE must be %s to mount this database (not %s)
ORA-00059: maximum number of DB_FILES exceeded
ORA-00060: deadlock detected while waiting for resource
ORA-00063: maximum number of LOG_FILES exceeded
ORA-00064: object is too large to allocate on this O/S (%s,%s)
ORA-00099: timed out while waiting for resource, potential PDML deadlock
ORA-00104: deadlock detected; all public servers blocked waiting for resources
ORA-00107: failed to connect to ORACLE listener process
ORA-00115: connection refused; dispatcher connection table is full
ORA-00125: connection refused; invalid presentation
ORA-00126: connection refused; invalid duplicity
ORA-00284: recovery session still in progress
ORA-00370: potential deadlock during kcbchange operation
ORA-00371: not enough shared pool memory
ORA-00376: file %s cannot be read at this time
ORA-00379: no free buffers available in buffer pool %s for block size %sK
ORA-00384: Insufficient memory to grow cache
ORA-00568: Maximum number of interrupt handlers exceeded
ORA-00579: osndnt: server received malformed connection request
ORA-00600: internal error code, arguments: [%s], [%s], [%s], [%s], [%s], [%s], [%s], [%s]
ORA-00603: ORACLE server session terminated by fatal error
ORA-01000: maximum open cursors exceeded
ORA-01012: not logged on
ORA-01014: ORACLE shutdown in progress
ORA-01019: unable to allocate memory in the user side
ORA-01031: insufficient privileges
ORA-01033: ORACLE initialization or shutdown in progress
ORA-01034: ORACLE not available
ORA-01035: ORACLE only available to users with RESTRICTED SESSION privilege
ORA-01037: maximum cursor memory exceeded
ORA-01046: cannot acquire space to extend context area
ORA-01073: fatal connection error: unrecognized call type
ORA-01089: immediate shutdown in progress - no operations are permitted
ORA-01090: shutdown in progress - connection is not permitted
ORA-01092: ORACLE instance terminated. Disconnection forced
ORA-01094: ALTER DATABASE CLOSE in progress. Connections not permitted
ORA-01109: database not open
ORA-01147: SYSTEM tablespace file %s is offline
ORA-01154: database busy. Open, close, mount, and dismount not allowed now
ORA-01155: the database is being opened, closed, mounted or dismounted
ORA-01219: database not open: queries allowed on fixed tables/views only
ORA-01237: cannot extend datafile %s
ORA-01456: may not perform insert/delete/update operation inside a READ ONLY transaction
ORA-01536: space quota exceeded for tablespace '%s'
ORA-01539: tablespace '%s' is not online
ORA-01542: tablespace '%s' is offline, cannot allocate space in it
ORA-01562: failed to extend rollback segment number %s
ORA-01573: shutting down instance, no further change allowed
ORA-01628: max # extents (%s) reached for rollback segment %s
ORA-01629: max # extents (%s) reached saving undo for tablespace %s

ORA-01630: max # extents (%s) reached in temp segment in tablespace %s
ORA-01631: max # extents (%s) reached in table %s.%s
ORA-01632: max # extents (%s) reached in index %s.%s
ORA-01650: unable to extend rollback segment %s by %s in tablespace %s
ORA-01651: unable to extend save undo segment by %s for tablespace %s
ORA-01652: unable to extend temp segment by %s in tablespace %s
ORA-01653: unable to extend table %s.%s by %s in tablespace %s
ORA-01654: unable to extend index %s.%s by %s in tablespace %s
ORA-01655: unable to extend cluster %s.%s by %s in tablespace %s
ORA-01656: max # extents (%s) reached in cluster %s.%s
ORA-01658: unable to create INITIAL extent for segment in tablespace %s
ORA-01659: unable to allocate MINEXTENTS beyond %s in tablespace %s
ORA-01680: unable to extend LOB segment by %s in tablespace %s
ORA-01681: max # extents (%s) reached in LOB segment in tablespace %s
ORA-01683: unable to extend index %s.%s partition %s by %s in tablespace %s
ORA-01684: max # extents (%s) reached in table %s.%s partition %s
ORA-01685: max # extents (%s) reached in index %s.%s partition %s
ORA-01686: max # files (%s) reached for the tablespace %s
ORA-01688: unable to extend table %s.%s partition %s by %s in tablespace %s
ORA-01691: unable to extend lob segment %s.%s by %s in tablespace %s
ORA-01692: unable to extend lob segment %s.%s partition %s by %s in tablespace %s
ORA-01693: max # extents (%s) reached in lob segment %s.%s
ORA-01694: max # extents (%s) reached in lob segment %s.%s partition %s
ORA-03113: end-of-file on communication channel
ORA-03114: not connected to ORACLE
ORA-03134: Connections to this server version are no longer supported.
ORA-03135: connection lost contact
ORA-03136: inbound connection timed out
ORA-03232: unable to allocate an extent of %s blocks from tablespace %s
ORA-03233: unable to extend table %s.%s subpartition %s by %s in tablespace %s
ORA-03234: unable to extend index %s.%s subpartition %s by %s in tablespace %s
ORA-03235: max # extents (%s) reached in table %s.%s subpartition %s
ORA-03236: max # extents (%s) reached in index %s.%s subpartition %s
ORA-03237: Initial Extent of specified size cannot be allocated
ORA-03238: unable to extend LOB segment %s.%s subpartition %s by %s in tablespace %s
ORA-03239: maxextents (%s) reached in LOB segment %s.%s subpartition %s
ORA-04020: deadlock detected while trying to lock object %s%s%s%s%s
ORA-06019: NETASY: invalid login (connect) string
ORA-06021: NETASY: connect failed
ORA-06030: NETDNT: connect failed, unrecognized node name
ORA-06031: NETDNT: connect failed, unrecognized object name
ORA-06032: NETDNT: connect failed, access control data rejected
ORA-06033: NETDNT: connect failed, partner rejected connection
ORA-06034: NETDNT: connect failed, partner exited unexpectedly
ORA-06035: NETDNT: connect failed, insufficient resources
ORA-06036: NETDNT: connect failed, no response from object
ORA-06037: NETDNT: connect failed, node unreachable
ORA-06039: NETDNT: connect failed
ORA-06040: NETDNT: invalid login (connect) string
ORA-06108: NETTCP: connect to host failed
ORA-06113: NETTCP: Too many connections
ORA-06114: NETTCP: SID lookup failure
ORA-06143: NETTCP: maximum connections exceeded
ORA-06315: IPA: Invalid connect string
ORA-06316: IPA: Invalid database SID
ORA-06317: IPA: Local maximum number of users exceeded
ORA-06318: IPA: Local maximum number of connections exceeded
ORA-06319: IPA: Remote maximum number of users exceeded
ORA-06320: IPA: Remote maximum number of connections exceeded
ORA-06404: NETCMN: invalid login (connect) string
ORA-06413: Connection not open.
ORA-10435: enable tracing of global enqueue service deadlock detection

ORA-10626: specify timeout for online index rebuild to wait for DML
ORA-10906: Unable to extend segment after insert direct load
ORA-12150: TNS:unable to send data
ORA-12151: TNS:received bad packet type from network layer
ORA-12152: TNS:unable to send break message
ORA-12153: TNS:not connected
ORA-12154: TNS:could not resolve service name
ORA-12155: TNS:received bad datatype in NSWMARKER packet
ORA-12156: TNS:tried to reset line from incorrect state
ORA-12157: TNS:internal network communication error
ORA-12158: TNS:could not initialize parameter subsystem
ORA-12159: TNS:trace file not writeable
ORA-12160: TNS:internal error: Bad error number
ORA-12161: TNS:internal error: partial data received
ORA-12162: TNS:service name is incorrectly specified
ORA-12163: TNS:connect descriptor is too long
ORA-12166: TNS:Client can not connect to HO agent.
ORA-12168: TNS:Unable to contact Directory Server.
ORA-12169: TNS:Net service name given as connect identifier is too long
ORA-12170: TNS:Connect timeout occurred
ORA-12171: TNS:could not resolve connect identifier: %s
ORA-12196: TNS:received an error from TNS
ORA-12197: TNS:keyword-value resolution error
ORA-12198: TNS:could not find path to destination
ORA-12200: TNS:could not allocate memory
ORA-12201: TNS:encountered too small a connection buffer
ORA-12202: TNS:internal navigation error
ORA-12203: TNS:unable to connect to destination
ORA-12204: TNS:received data refused from an application
ORA-12205: TNS:could not get failed addresses
ORA-12206: TNS:received a TNS error during navigation
ORA-12207: TNS:unable to perform navigation
ORA-12208: TNS:could not find the TNSNAV.ORA file
ORA-12209: TNS:encountered uninitialized global
ORA-12210: TNS:error in finding Navigator data
ORA-12211: TNS:needs PREFERRED_CMANAGERS entry in TNSNAV.ORA
ORA-12212: TNS:incomplete PREFERRED_CMANAGERS binding in TNSNAV.ORA
ORA-12213: TNS:incomplete PREFERRED_CMANAGERS binding in TNSNAV.ORA
ORA-12214: TNS:missing local communities entry in TNSNAV.ORA
ORA-12215: TNS:poorly formed PREFERRED_NAVIGATORS Addresses in TNSNAV.ORA
ORA-12216: TNS:poorly formed PREFERRED_CMANAGERS addresses in TNSNAV.ORA
ORA-12217: TNS:could not contact PREFERRED_CMANAGERS in TNSNAV.ORA
ORA-12218: TNS:unacceptable network configuration data
ORA-12219: TNS:missing community name from address in ADDRESS_LIST
ORA-12221: TNS:illegal ADDRESS parameters
ORA-12222: TNS:no such protocol adapter
ORA-12223: TNS:internal limit restriction exceeded
ORA-12224: TNS:no listener
ORA-12225: TNS:destination host unreachable
ORA-12226: TNS:operating system resource quota exceeded
ORA-12227: TNS:syntax error
ORA-12228: TNS:protocol adapter not loadable
ORA-12229: TNS:Interchange has no more free connections
ORA-12230: TNS:Severe Network error occurred in making this connection
ORA-12231: TNS:No connection possible to destination
ORA-12232: TNS:No path available to destination
ORA-12233: TNS:Failure to accept a connection
ORA-12235: TNS:Failure to redirect to destination
ORA-12236: TNS:protocol adapter not loaded
ORA-12316: syntax error in database link's connect string
ORA-12326: database %s is closing immediately; no operations are permitted
ORA-12329: database %s is closed; no operations are permitted

ORA-12500: TNS:listener failed to start a dedicated server process
ORA-12501: TNS:listener failed to spawn process
ORA-12502: TNS:listener received no CONNECT_DATA from client
ORA-12504: TNS:listener was not given the SID in CONNECT_DATA
ORA-12505: TNS:listener could not resolve SID given in connect descriptor
ORA-12506: TNS:listener was not given the ALIAS in CONNECT_DATA
ORA-12507: TNS:listener could not resolve ALIAS given
ORA-12508: TNS:listener could not resolve the COMMAND given
ORA-12509: TNS:listener failed to redirect client to service handler
ORA-12510: TNS:database temporarily lacks resources to handle the request
ORA-12511: TNS:service handler found but it is not accepting connections
ORA-12512: TNS:service handler found but it has not registered a
redirect address
ORA-12513: TNS:service handler found but it has registered for a
different protocol
ORA-12514: TNS:listener could not resolve SERVICE_NAME given in
connect descriptor
ORA-12515: TNS:listener could not find a handler for this presentation
ORA-12516: TNS:listener could not find available handler with matching
protocol stack
ORA-12517: TNS:listener could not find service handler supporting direct handoff
ORA-12518: TNS:listener could not hand off client connection
ORA-12519: TNS:no appropriate service handler found
ORA-12520: TNS:listener could not find available handler for requested
type of server
ORA-12521: TNS:listener could not resolve INSTANCE_NAME given in
connect descriptor
ORA-12522: TNS:listener could not find available instance with given
INSTANCE_ROLE
ORA-12523: TNS:listener could not find instance appropriate for the
client connection
ORA-12524: TNS:listener could not resolve HANDLER_NAME given in
connect descriptor
ORA-12525: TNS:listener has not received client's request in time allowed
ORA-12526: TNS:listener: all appropriate instances are in restricted mode
ORA-12527: TNS:listener: all instances are in restricted mode or blocking
new connections
ORA-12528: TNS:listener: all appropriate instances are blocking new connections
ORA-12529: TNS:connect request rejected based on current filtering rules
ORA-12531: TNS:cannot allocate memory
ORA-12532: TNS:invalid argument
ORA-12533: TNS:illegal ADDRESS parameters
ORA-12534: TNS:operation not supported
ORA-12535: TNS:operation timed out
ORA-12536: TNS:operation would block
ORA-12537: TNS:connection closed
ORA-12538: TNS:no such protocol adapter
ORA-12539: TNS:buffer over- or under-flow
ORA-12540: TNS:internal limit restriction exceeded
ORA-12541: TNS:no listener
ORA-12542: TNS:address already in use
ORA-12543: TNS:destination host unreachable
ORA-12544: TNS:contexts have different wait/test functions
ORA-12545: Connect failed because target host or object does not exist
ORA-12546: TNS:permission denied
ORA-12547: TNS:lost contact
ORA-12549: TNS:operating system resource quota exceeded
ORA-12550: TNS:syntax error
ORA-12551: TNS:missing keyword
ORA-12552: TNS:operation was interrupted
ORA-12554: TNS:current operation is still in progress
ORA-12555: TNS:permission denied

```
ORA-12556: TNS:no caller
ORA-12557: TNS:protocol adapter not loadable
ORA-12558: TNS:protocol adapter not loaded
ORA-12560: TNS:protocol adapter error
ORA-12561: TNS:unknown error
ORA-12562: TNS:bad global handle
ORA-12564: TNS:connection refused
ORA-12566: TNS:protocol error
ORA-12569: TNS:packet checksum failure
ORA-12570: TNS:packet reader failure
ORA-12571: TNS:packet writer failure
ORA-12574: TNS:redirection denied
ORA-12582: TNS:invalid operation
ORA-12583: TNS:no reader
ORA-12585: TNS:data truncation
ORA-12589: TNS:connection not bequeathable
ORA-12590: TNS:no I/O buffer
ORA-12591: TNS:event signal failure
ORA-12592: TNS:bad packet
ORA-12593: TNS:no registered connection
ORA-12595: TNS:no confirmation
ORA-12596: TNS:internal inconsistency
ORA-12600: TNS: string open failed
ORA-12602: TNS: Connection Pooling limit reached
ORA-12606: TNS: Application timeout occurred
ORA-12607: TNS: Connect timeout occurred
ORA-12608: TNS: Send timeout occurred
ORA-12609: TNS: Receive timeout occurred
ORA-12612: TNS:connection is busy
ORA-12615: TNS:preempt error
ORA-12623: TNS:operation is illegal in this state
ORA-12624: TNS:connection is already registered
ORA-12636: Packet send failed
ORA-12637: Packet receive failed
ORA-12829: Deadlock - itls occupied by siblings at block %s of file %s
ORA-12993: tablespace '%s' is offline, cannot drop column
ORA-14117: partition resides in offlined tablespace
ORA-14268: subpartition '%s' of the partition resides in offlined tablespace
ORA-16000: database open for read-only access
ORA-16003: standby database is restricted to read-only access
ORA-16403: shutdown in progress - remote connection is not permitted
ORA-16724: the intended state for resource has been set to OFFLINE
ORA-16903: Unable to connect to database
ORA-16914: Missing connect string. Try \"help\"
ORA-18014: deadlock detected while waiting for resource %
ORA-21521: exceeded maximum number of connections in OCI (object mode only)
ORA-21522: attempted to use an invalid connection in OCI (object mode only)
ORA-23317: a communication failure has occurred
ORA-24401: cannot open further connections
ORA-24418: Cannot open further sessions.
ORA-24778: cannot open connections
ORA-25400: must replay fetch
ORA-25401: can not continue fetches
ORA-25402: transaction must roll back
ORA-25403: could not reconnect
ORA-25405: transaction status unknown
ORA-25407: connection terminated
ORA-25408: can not safely replay call
ORA-25409: failover happened during the network operation, cannot continue
ORA-25425: connection lost during rollback
ORA-29306: datafile %s is not online
ORA-30006: resource busy; acquire with WAIT timeout expired
```

ORA-30036: unable to extend segment by %s in undo tablespace '%s'
ORA-30040: Undo tablespace is offline
ORA-31443: deadlock detected while acquiring lock on %s
ORA-37013: (XSACQUIRE_DEADLOCK) Cannot wait to acquire object %j since doing so
would cause a deadlock.
ORA-44317: database open read-only