

Terracotta Management and Monitoring

Version 10.7

October 2020

This document applies to Terracotta 10.7 and to all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2010-2022 Software AG, Darmstadt, Germany and/or Software AG USA, Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <https://softwareag.com/licenses/>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <https://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <https://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

Document ID: TC-TMC-UG-107-20221101

Table of Contents

About This Documentation	5
Online Information and Support.....	6
Data Protection.....	6
1 Getting Started with the Terracotta Management Console	7
2 Using the Terracotta Management Console	11
3 Using the Ehcache Tab	23
4 Using the TCStore Tab	33
5 Using the Server Tab	49
6 Using the Events Tab	53
7 Prometheus Integration	55
8 Performance Considerations	61
9 Migrating the Terracotta Management Console	63

About This Documentation

- Online Information and Support 6
- Data Protection 6

Online Information and Support

Software AG Documentation Website

You can find documentation on the Software AG Documentation website at <http://documentation.softwareag.com>.

Software AG Empower Product Support Website

If you do not yet have an account for Empower, send an email to empower@softwareag.com with your name, company, and company email address and request an account.

Once you have an account, you can open Support Incidents online via the eService section of Empower at <https://empower.softwareag.com/>.

You can find product information on the Software AG Empower Product Support website at <https://empower.softwareag.com>.

To submit feature/enhancement requests, get information about product availability, and download products, go to [Products](#).

To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the [Knowledge Center](#).

If you have any questions, you can find a local or toll-free number for your country in our Global Support Contact Directory at https://empower.softwareag.com/public_directory.aspx and give us a call.

Software AG TECHcommunity

You can find documentation and other technical information on the Software AG TECHcommunity website at <http://techcommunity.softwareag.com>. You can:

- Access product documentation, if you have TECHcommunity credentials. If you do not, you will need to register and specify "Documentation" as an area of interest.
- Access articles, code samples, demos, and tutorials.
- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.
- Link to external websites that discuss open standards and web technology.

Data Protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR). Where applicable, appropriate steps are documented in the respective administration documentation.

1 Getting Started with the Terracotta Management Console

About the Terracotta Management Console

This document describes the Terracotta Management Console (TMC), which is a browser-based application served by the Terracotta Management Server (TMS). The TMC provides a complete view of your Terracotta Server Array (TSA) and connected clients.

With the TMC you can:

- observe the cluster topology and categories of connected clients
- view statistics
- clear cache contents
- and more

Installing the TMS

The TMS is a standard Java web application and bundled web container. It gathers information from throughout the cluster and serves it to the TMC for display in your browser. The TMS is located in the installed Terracotta kit under the `tools/management` directory.

Configuring the TMS

Certain aspects of the TMS can be customized via the properties file located in the installed Terracotta kit at `tools/management/conf/tmc.properties`. While that file contains many more properties, those that follow are the most likely to be useful to the TMS administrator:

```
server.port=9480
tms.defaultUrl=terracotta://localhost:9410
spring.pid.file=${tmc.home:build/tmc}/run/tmc.pid
tms.storageFolder=target/tmc/data
tms.offheapSizeMB=1024
tms.persistenceMode=HYBRID
tms.offheapMonitoringFrequencyMs=30000
tms.offheapThreshold=80
tms.statisticsMaxAgeMinutes=1500
tms.statisticsScavengerFrequencyMs=30000
tms.eventLogMaxRecords=5000
```

```
tms.eventLogScavengerFrequencyMs=60000
```

Security

Please see the section *TMS Security* in the *Terracotta Server Administration Guide* for information on how to configure TMS security.

In essence, you need to configure the following properties in the `tmc.properties` file:

```
tms.security.root.directory
tms.security.audit.directory
tms.security.https.enabled
tms.security.authentication.scheme
tms.security.authorization.scheme
tms.security.root.directory.connection.default
```

and provide the correct files in the security root directory.

Starting and Stopping the TMS

To start the TMS, execute the script `start.bat` (on Windows) or `start.sh` (on UNIX-based systems), located in `tools/management/bin` under the installed Terracotta kit.

To terminate the TMS, use the associated `stop.bat` or `stop.sh` script, or platform-provided process management tools.

Adding Manageability to your Ehcache CacheManager

In order to get the fullest manageability and monitorability you must configure your CacheManager to make use of a ManagementRegistryService:

```
import org.ehcache.config.builders.CacheManagerBuilder;
import org.ehcache.management.registry.DefaultManagementRegistryConfiguration;
CacheManager cacheManager = CacheManagerBuilder.newCacheManagerBuilder()
    .using(new DefaultManagementRegistryConfiguration()
        .setCacheManagerAlias("MyCacheManager")
        .addTags("my-client-tag", "another-client-tag"))
    .build(true);
```

Configuring the management registry will activate it and will allow various capabilities to be exposed via the TMS REST interface with the alias and tags you provide. We recommend you to assign the same alias to equivalent instances of your CacheManager, across different clients so as to support statistics aggregation.

If you *do not* make use of a management registry, your CacheManager will still appear in the TMC but most management and monitoring features will not operate.

Tip:

If you want to exclude a cache manager from the TMC UI, you can add the following tag: `tmc-excluded`.

For related information, see the section *Ehcache API Developer Guide > Management and Monitoring with Ehcache*.

Adding Manageability to your TCStore Dataset

Unlike for an Ehcache CacheManager, there is no need to explicitly add manageability to your TCStore Dataset.

Still, if you are using a clustered dataset manager, it is good practice to set an alias to your dataset and some tags:

```
DatasetManager datasetManager = DatasetManager.clustered([...])
    .withClientAlias("myDsManager")
    .withClientTags("node-1", "webapp-2", "testing")
    .build();
```

Tip:

If you want to exclude a dataset manager from the TMC UI, you can add the following tag: `tmc-excluded`.

Connecting to the TMC

After starting the TMS, open a browser and visit `http://localhost:9480`. The TMC will load and present you with the Home Page, where persistent connections to your clusters can be created, viewed and managed.

Screenshot: TMC Home Page, Create New Connection



[+ Create New Connection](#)

2 Using the Terracotta Management Console

The TMC Home Page

The TMC home page is where you:

- create/delete persistent connections to your cluster(s)
- optionally modify your connection properties
- view
 - the status of the servers that make up your cluster
 - the various categories of clients making use of your cluster
 - the server entities that are contained by the cluster and to which clients connect
- drill-down/jump to various presentations such as statistics and monitoring relating to those servers, clients, and entities
- take actions, such as clearing the contents of a cache

Tip:

TMC requires browser cookies to work. Please make sure your browser doesn't block reading cookies (some browser plugins can block reading cookies). If you see a related error message, you can try another browser, or use incognito mode, or disable browser plugins.

For information about starting and stopping Terracotta servers, refer to the section *Starting and Stopping the Terracotta Server* in the *Terracotta Server Administration Guide*.

For information about creating Terracotta clusters, refer to the section *The Cluster Tool* in the *Terracotta Server Administration Guide*.

Connections and Global Settings

To create a persistent connection to a particular operational cluster:

1. Click *Create New Connection*
2. In the *Connection URL* input area enter a URL addressing at least a single, running member of the fully configured Terracotta Server Array (TSA):

Screenshot: TMC Home Page. Connection URL to at least one running member of the TSA



Terracotta Server URL

```
terracotta://<server-host>:<listen-port>[,<server-host>:<listen-port>]*
```

If the TMS is able to connect to the specified server, it will request the complete TSA topology, persisting the addresses of each server in its database. This means that, in the future, the TMS will be able to connect to the TSA, even if the originally specified server should be unreachable, as long as at least one member of the TSA is running.

If you attempt to connect to a running server that is not part of a *configured* cluster, the TMS will return an error indicating that no license was discovered.

Note:

It is a best-practice to specify the addresses of all the members of a particular TSA stripe in the connection URL to support high-availability. Simply comma-separate each server's address (host:port).

The TMS will attempt to connect to each server, in turn, until a successful connection is established.

3. Enter the name you would like associated with this connection, for example "MyCluster", then click **Next** to create the connection.

Screenshot: TMC Home Page, Connection Name



Connection Name:

This connection name will be used to disambiguate multiple connections in the UI and when communicating with the TMS REST interface.

4. Upon completion of creating the new connection, a new *connection region* is shown on the Home Page, displaying the current state of your cluster.

Screenshot: TMC Home Page, State of the TSA after initial entries.



[+ Create New Connection](#)

MyCluster

Ehcache TCStore Server Event Log Versions Export Diagnostic Data ↓ ⏻ ✎ 🗑️

▶ Server Array **stripes(2): active(2) passive(2) unreachable(0)**

▶ **Connected Clients (6)**

▶ Ehcache Server Entities **total(1) in-use(1)**

▶ TCStore Server Entities **total(5) in-use(5)**

▶ Websessions Entities **total(4) in-use(4)**

Copyright © - Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

software AG

The connection to MyCluster in the example above shows a TSA comprised of 2 stripes, each containing a pair of Active-Passive servers, 4 currently connected clients, a single Ehcache server entity and a single TCStore server entity. Further detail can be exposed by drilling down into the display hierarchy.

Note:

For Terracotta Ehcache product users, the **TCStore** dropdown is deactivated.

Modifying a Connection's Properties

If you wish to modify your new connection's properties, use the *Edit this cluster* () icon on the connection's header area.

Screenshot: TMC Cluster Properties Editor (Overlay)



Edit: MyCluster [Close]

Connection Timeout: 10 [Up] [Down]

Read Timeout: 10 [Up] [Down]

Collector Interval: 30 [Up] [Down]


[Apply]

Connection Timeout	Timeout value, in seconds, when connecting to a cluster the first time or when reconnecting. Default is 10 seconds.
Read Timeout	Timeout value, in seconds, for all cluster communications. Default is 10 seconds.
Collector Interval	Collector interval for statistics. Default is 30 seconds. Minimum is 10 seconds and maximum is 60 seconds.

Deleting a Connection

To delete an existing connection use the *Delete this cluster* () icon on the connection header area.

Screenshot: TMC Connection Area Header



▲ **MyCluster**

Ehcache [v] TCStore [v] Server [v] Event Log Versions Export Diagnostic Data

Using the Configured Connections

The TMC home page shows each of your configured connections in its own collapsible region of the display that presents the totality of the cluster, allowing you to drill-down to different levels of hierarchy then jump to various detail views, such as for statistics or an entity's server-side resource usage.

Each configured connection region is comprised of several high-level facets which will now be described.

Buttons in the connection region header

The header of the connection region for a selected cluster contains the following selectable buttons:

- **Ehcache**
- **TCStore**
- **Server**
- **Events**
- **Export Diagnostic Data**

If you select the **Ehcache**, **TCStore**, **Server** or **Events** button, the **Detail** page will be displayed, with the appropriate tab (**Ehcache**, **TCStore**, **Server** or **Events**) already selected.

Each of these three buttons contains a dropdown menu. If you select an item from any of these dropdown menus, the **Detail** page will be displayed and the details for the selected dropdown entry will be in focus.


For details of the **Ehcache**, **TCStore**, **Server** and **Events** tabs of the **Detail** page, see the sections “Using the Ehcache Tab” on page 23, “Using the TCStore Tab” on page 33, “Using the Server Tab” on page 49 and “Using the Events Tab” on page 53.

For information about the button **Export Diagnostic Data**, see the following topic.

Export Diagnostic Data (Cluster level)

The **Export Diagnostic Data** button at the cluster level provides diagnostic information for every server in the cluster.

Screenshot: TMC Connection Area Header, Highlighting button for Export Diagnostic Data



The cluster level **Export Diagnostic Data** button exports the following diagnostic data as a zip file.

- **Event Log** - a csv file named `EventLog.csv`, containing cluster wide events. This can also be exported from the **Events** panel.

- Versions - a text file named `versions.txt`, containing build version information for all nodes in the cluster and their connected clients.
- For each server it also downloads seven diagnostic data files, which are explained in the next section, *Export Diagnostic Data (Server level)*.

The zip file is named using the format: `diagnostic-connectionName-dateTime.zip`

- e.g. `diagnostic-myClusterConnection-20170705221353.zip`

and the layout of the zip file is:

- one folder for each stripe, containing one folder for each server within the stripe
- each server folder contains the seven files referenced in the next section, *Export Diagnostic Data (Server level)*
- in the root of the zip file there is the cluster wide event log, `EventLog.csv`, and a `versions.txt` file which lists the versions of all nodes in the cluster and their connected clients.

Server Array

The *Terracotta Server Array (TSA)* is a collection of groups of servers, known as *stripes*. Servers within a stripe work together to provide *High-Availability*. If the *Active* server should fail, any one of the remaining *Passive* servers takes over as active. The *Active* server serves to (1) handle requests from clients to entities it contains and (2) to relay those client requests to each of the *Passive* servers. Servers in different stripes do not interact.

The TMC presents the current state of the server array, indicating the roles of each server. The following shows a server array consisting of two stripes, each containing two members, one *active* and one *passive*.

Screenshot: Example of a TMC Terracotta Server Array, two stripes, each having two members, one active, one passive.

The screenshot displays the following structure:

- Server Array **stripes(2): active(2) passive(2) unreachable(0)**
 - stripe[0] **active(1) passive(1) unreachable(0)**
 - terracotta-1-0.stripe-1:9410 (terracotta-1...) View Diagnostic Details Export Diagnostic Data Terminate Server (Active)
 - terracotta-1-1.stripe-1:9410 (terracotta-1...) View Diagnostic Details Export Diagnostic Data Terminate Server (Passive)
 - stripe[1] **active(1) passive(1) unreachable(0)**
 - terracotta-2-0.stripe-2:9410 (terracotta-2...) View Diagnostic Details Export Diagnostic Data Terminate Server (Active)
 - terracotta-2-1.stripe-2:9410 (terracotta-2...) View Diagnostic Details Export Diagnostic Data Terminate Server (Passive)

The possible server states are:


Icon	Server state	Description
	STARTING	server is starting
	UNINITIALIZED	server has started and is ready for election
	SYNCHRONIZING	server is synchronizing its data with the current active server
	PASSIVE	server is passive and ready for replicatio
	ACTIVE	server is active and ready to accept client
	ACTIVE_RECONNECTING	server is active but waits for previously known clients to rejoin before accepting new clients
	ACTIVE_SUSPENDED	server is active but blocked in the election process (consistency mode)
	PASSIVE_SUSPENDED	server is passive but blocked in the election process (consistency mode)
	START_SUSPENDED	server is starting but blocked in the election process (consistency mode)
	UNREACHABLE	server is unreachable from TM
	UNKNOWN	default server state shown in TMC before TMC is able to access the server stat

Diagnostic information for the members of your TSA can be downloaded as an archive file or viewed directly.

View Diagnostic Details (Server level)

The **View Diagnostic Details** dropdown at the server level lets you view directly the diagnostic information for that particular server. There is one diagnostic data dropdown per server and it is always positioned to the immediate right of the server name.

Screenshot: Example of a TMC Terracotta Server Array, dropdown highlighted in action.

**Note:**

Please don't confuse the stripe count with the stripe name. In this example there are two stripes, indicated by **stripes(2)**, and they are named `stripe[0]` and `stripe[1]`.

Export Diagnostic Data (Server level)

The **Export Diagnostic Data** button at the server level lets you download diagnostic information for that particular server in an archive file, similar in format to that provided at the cluster level.

The dropdown provides access to the seven diagnostic data files below. The file contents can be viewed/exported individually by selecting that option from the dropdown or by selecting the top level **Export Diagnostic Data** option which exports a zip file containing all of them.

1. Environment - shows a list of all the environment variables.
2. TC Properties - provides a list of all the TC config properties.
3. Process Arguments - displays all the command line arguments submitted for the process.
4. Cluster Configuration - shows the Terracotta configuration file.
5. Cluster State - displays information on the current cluster state.
6. Thread Dumps - exports the thread dump as a txt file.
7. Component Versions - displays a listing of build version information for all servers and their connected clients.

Tip:

What is a thread dump? A Java thread dump is a way of finding out what every thread in the JVM is doing at a particular point in time. This is especially useful if your Java application sometimes seems to hang when running under load, as an analysis of the dump will show where the threads are stuck.

The zip file is named using the format:

`diagnostic-connectionName-stripeName-serverName-dateTime.zip`

- e.g. `diagnostic-myClusterConnection-stripe[0]-testServer0-20170705221356.zip`

and the layout of the zip file is one folder, named after the server name, containing the 6 data files above.

Connected Clients

In the Terracotta Platform, a *client* is an application end-point. In your application, the clustered CacheManager that is configured and initialized is a Terracotta client. Each client maintains a connection to the active server in each stripe. In general, anything that connects to a server is considered a client.

For more information, see section *Terracotta Server Administration Guide > Clients in a Cluster*.

Each client has a *Client Identifier* that serves to uniquely identify that client. The form of the identifier is:

```
<pid>@<ip-addr>:<client-type>:<server-entity-name>:<uuid>
```


where:

<pid>	the <i>process identifier</i> of the Java Virtual Machine hosting this client
<ip-addr>	the <i>IP address</i> of the machine hosting this JVM
<client-type>	[Ehcache Store Unknown]
<server-entity-name>	the name of the server-side entity the client is connected to
<uuid>	a unique identifier that serves to disambiguate clients in the same JVM, accessing the same server entity

■ Ehcache Clients

The following shows two caching clients.

Screenshot: Example Connected Clients: Two caching clients.



The input field located next to *Connected Clients* serves to show only those clients whose identifier contains the entered value. In the example above, entering 81390 (the *process identifier* of the 2nd client) would filter out the first client.

Filtering Rules


- accepts a space-separated list of terms to match (prefix with ! to negate)
- terms can apply to any components of the *Client Identifier* or any supplied tags

- negated terms must not match any of the above
- clear the input field content or click the toggle button () to show all clients

Under each identifier is an optional set of user-defined tags, specified in the CacheManager's configuration.


Expanding a *caching client* exposes the CacheManager's alias, which is defined in the configuration. The dropdown lets you jump to various detail views, pre-selecting this client and CacheManager in those views. View the CacheManager's configuration by selecting the **configuration** () icon. The cache entries can be cleared by selecting the **clear cache** () icon.

Screenshot: TMC Cache Manager's Configuration, elements collapsed



Expanding the cache shows important configuration elements related to that cache.

Screenshot: TMC Caching ClientExpanded



■ TCStore Clients


Within the connected clients section there is also a subsection for all the dataset clients.

Expanding the **TCStore Clients** section provides a list of all the connected clients, ordered by client identifiers.

Further expanding a client identifier shows all datasets associated with that particular client and for each dataset there is a dropdown to navigate quickly to its overview or chart statistics.

Lastly when expanding each dataset there is also a list of all its dataset instances

Screenshot: Expanding a Client Identifier



Ehcache Entities

The Terracotta Platform consists of client-side programmatic (API) artifacts and server-side entities that work together to provide highly-available, performant, distributed data access. In

your application you configure a CacheManager, both on the client- and server-side. The client-side configuration relates to such things as the maximum OnHeap size-in-bytes for a particular cache. The server-side configuration relates to the remote storage tier that is used to store your cache entries.

Each of your *Ehcache Clients* communicates with its remote storage tier when executing normal cache operations, such as putting an entry into the cache.

Ehcache clients configured to use the same remote storage tier are effectively sharing access to the same cache data.

TCStore Entities

A TCStore client is essentially a connection to a clustered or embedded Dataset.

A single client instance can *fetch* an arbitrary number of *handles* to the underlying Dataset, referred to as a Dataset instance. Operations statistics are maintained on a per-instance basis.

3 Using the Ehcache Tab

The **Ehcache** tab of the **Detail** page contains detailed presentations relating to your caching clients. The **Detail** page is selectable via buttons in the home page.

You can view:

- Overview statistics concerning this moment in time, in grid form. Select this view by clicking the icon .
- Historical statistics, in chart form. Select this view by clicking the icon .
- Cache size information. Select this view by clicking the icon .

Your caching client automatically gathers and periodically sends to the TMS low-level statistics concerning counters and sizing. From those low-level statistics the TMS can synthesize a variety of derived statistics, such as rates (HitRate) and ratios (HitRatio).

The following categories of statistics are available:

Category	Description
Cache	encompasses all the following caching tiers
OnHeap	the JVM heap tier
OffHeap	the JVM OffHeap (direct memory) tier
Disk	the disk tier
Clustered	the cluster tier

The following raw counter values, as well as associated rate of change, are available for each category listed above:

Counter	Description
<Category>:PutCount	number of times your application has put a new entry into the cache
<Category>:UpdateCount	number of times your application has put a new entry into the cache that replaced an existing entry (same key)

Counter	Description
<Category>.RemovalCount	number of times your application explicitly removed an entry from the cache
<Category>.ExpirationCount	number of entries removed from the cache due to an expiration policy (time-to-live, etc.)
<Category>.EvictionCount	number of entries removed from the cache due to space constraints
<Category>.HitCount	number of gets that returned an existing entry
<Category>.HitRatio	ratio of hits to gets
<Category>.MissCount	number of gets that did not return an entry
<Category>.MissRatio	ratio of misses to gets

The following category-specific raw counters are also available:

Counter	Description
Disk:MappingCount	number of entries stored on disk
Disk:AllocatedByteSize	number of bytes allocated for storage. AvailableByteSize = AllocatedByteSize - OccupiedByteSize
Disk:OccupiedByteSize	total size of all entries stored on disk
OffHeap:MappingCount	number of entries stored in direct memory
OffHeap:AllocatedByteSize	number of bytes allocated off heap in memory. AvailableByteSize = AllocatedByteSize - OccupiedByteSize
OffHeap:OccupiedByteSize	total size of all entries stored in direct memory
OnHeap:MappingCount	number of entries that store in regular heap memory
OnHeap:OccupiedByteSize	total size of all entries store in regular heap memory

The following latencies are also available:

Counter	Description
Cache:GetHitLatency	latencies of get operations leading to a cache hit. 4 different percentiles are returned: median, 95th, 99th and maximum. <ul style="list-style-type: none"> ■ Cache:GetHitLatency#50 ■ Cache:GetHitLatency#95

Counter	Description
	<ul style="list-style-type: none"> ■ Cache:GetHitLatency#99 ■ Cache:GetHitLatency#100
Cache:GetMissLatency	<p>latencies of get operations leading to a cache miss. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Cache:GetMissLatency#50 ■ Cache:GetMissLatency#95 ■ Cache:GetMissLatency#99 ■ Cache:GetMissLatency#100
Cache:PutLatency	<p>latencies of successful put operations. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Cache:PutLatency#50 ■ Cache:PutLatency#95 ■ Cache:PutLatency#99 ■ Cache:PutLatency#100
Cache:GetRemoveLatency	<p>latencies of successful remove operations. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Cache:RemoveLatency#50 ■ Cache:RemoveLatency#95 ■ Cache:RemoveLatency#99 ■ Cache:RemoveLatency#100

Note:

Ratios and Rates

The TMS gathers raw counters from your application and the servers. Rates of change (rates) and ratios are synthesized based on windowed sample aggregation. The aggregation method depends on the statistic type. For counters and latencies, the maximum value observed during an interval of time is displayed. The TMC requests a different number of samples depending on the chart's time-frame, which will determine the length of the window to use for rates and ratio calculation. The units for rates are always operations per second. Ratios are unitless.

Overview Panel

The Overview Panel lets you view statistics concerning what is happening right now (or just a while ago) with your caching client, broken out by cache. By default the statistics are aggregated

across all clients using the same CacheManager. Any particular caching client can also be selected for viewing.

Screenshot: TMC Caching Overview, listing caches and statistics.

CacheName	Cache:HitCount	Cache:HitRate	Cache:HitRatio	Cache:MissCount	Cache:MissRate	Cache:MissRatio
dedicatedcache -> dedicatedcache	2.8K	4.07	100 %	5.0	0.00	0 %
more-cache-4 -> more-cache-4	2.7K	3.90	100 %	0	0.00	0 %
more-cache-5 -> more-cache-5	2.7K	4.00	100 %	0	0.00	0 %
more-cache-6 -> more-cache-6	2.8K	4.03	100 %	0	0.00	0 %
shared-cache-1 -> shared-cache-1	2.7K	4.10	100 %	0	0.00	0 %
shared-cache-2 -> shared-cache-2	2.8K	3.90	100 %	3.0	0.00	0 %

Use the **CacheManagers** dropdown to select the aliased CacheManager for which to show statistics. The entries in this dropdown are of the form *CacheManager Alias > Storage Handle* due to the fact that different client-side CacheManagers can be configured to use the same storage on the TSA. Furthermore, CacheManagers that are otherwise identical can be configured with different aliases. We show the mapping here so you have a chance to disambiguate CacheManagers should you choose to configure them in this way.

Use the **Clients** dropdown to select a particular caching client for viewing. By default, statistics are aggregated across all clients of the selected CacheManager.

Additional Grid Features


- Temporarily re-arrange grid columns via drag-and-drop
- Sort on columns or filter on CacheName
- Scroll horizontally if the grid columns overflow the available space

Use the **Export** () icon to download a spreadsheet of the currently displayed values.

Use the **Filter cache statistics** () icon to select which statistics to display. The set of statistics to view defaults to **Cache**-level statistics that are applicable to all use-cases.

You can also filter statistics by name, using space-separated terms.

Screenshot: Overlay, choosing statistics



Charts Panel

The **Charts** Panel lets you view statistics over the past 5 minutes. The TMC requests 30 samples be returned. Unlike the **Overview** Panel, which shows all the caches contained by the selected CacheManager, the **Charts** Panel shows statistics for particular caches.

Screenshot: Charts Panel



Use the **CacheManagers** dropdown to select the aliased CacheManager for which to show statistics. The entries in this dropdown are of the form **CacheManager Alias > Storage Handle** due to the fact that different client-side CacheManagers can be configured to use the same storage on the TSA. Furthermore, CacheManagers that are otherwise identical can be configured with different aliases. We show the mapping here so you have a chance to disambiguate CacheManagers should you choose to configure them in this way.

Use the **Clients** dropdown to select a particular caching client for viewing. By default, statistics are aggregated across all clients of the selected CacheManager.

Use the **Caches** dropdown to select a particular cache for viewing. By default, the first cache listed is selected.



Use the slider (2 3 4) to set how many columns of charts you would like displayed.

Use the events checkbox ( **Events**) to choose whether or not you want to display clients and cache cleared events on the charts. This causes additional dotted lines to appear on the charts, showing the time when the following events occurred: "EC" = Ehcache Cleared, "CJ" = Client Joined, "CL" = Client Left.

Use the **Filter cache statistics** () icon to select which statistics to display. The set of statistics to view defaults to **Cache**-level statistics that are applicable to all use-cases.

Use the **Take a snapshot of all charts** () icon to download a single PNG file containing the current values of all displayed charts.

Directly under each individual chart, use the **Export to PDF** () icon or the **Export to PNG image** () icon to download the chart in the selected format.

Additional Chart Features

- Rearrange charts via drag-and-drop; the order is preserved until the browser is closed or its cache is manually cleared. If new statistics are selected, they will appear after the charts, and then you can rearrange them.
- Double-click a chart to get an enlarged snapshot
- Check the box named "Events" to display "Client Joined" , "Client Left" and "Ehcache Cleared" events on your charts


Sizing Panel

The Sizing Panel shows how much space is being used by your CacheManagers and caches across the different local storage tiers you've configured. Sizing information related to the *Clustered Storage Tier* can be accessed via the **Remote** buttons.

The **Sizing** Panel is composed of two sections: *CacheManager Level Sizing* and *Cache Level Sizing*.

- **CacheManager Level Sizing**

Screenshot: Sizing Panel



Shown is a pie chart for each local storage tier you've configured for use by your caches, displaying the amount of the available storage that is being used.


Use the **CacheManagers** dropdown to select the aliased CacheManager for which to display sizing information.

Use the **Clients** dropdown to select a particular caching client for which to display sizing information. By default the first client listed is selected.

Use the **Remote** button to navigate to the Resource Usage panel to view how the server-side caching entity is making use of the server's configured resources, such as OffHeap storage. See the section "[Resource Usage Panel](#)" on page 49 for related information.

■ Cache Level Sizing

Screenshot: Cache Level Sizing



Use the **View By Tier** dropdown to view tier-specific sizing information for all caches contained by the selected CacheManager in grid form.

Note:

Byte-sizing Limitations

Cache tiers that are sized by *entries* cannot provide *Occupied Size* or *Average Entry Size* and are displayed as N/A for *Not Applicable*.

Use the **View By Cache** dropdown to view the tier usage breakdown for a particular cache, in pie chart form.

4 Using the TCStore Tab

The **TCStore** tab of the **Detail** page opens a view containing detailed statistics for TCStore dataset operations. The **Detail** page is selectable via buttons in the home page.

The following raw counter values, as well as associated rate of change, are available:

Counter	Description
Add:AlreadyExists	count of adds that failed due to a record with the specified key already existing in the dataset
Add:Failure	count of unsuccessful additions to a dataset
Add:Success	count of successful additions to a dataset
Delete:Failure	count of unsuccessful deletions from a dataset
Delete:NotFound	count of deletes that failed due to a record with the specified key not existing in the dataset
Delete:Success	count of successful deletions from a dataset
Get:Failure	count of unsuccessful gets from a dataset
Get:NotFound	count of failed gets that were due to no record with the specified key existing in the dataset
Get:Success	count of successful gets from a dataset
Update:Failure	count of unsuccessful attempts to update a record in the dataset
Update:NotFound	count of update failures due to no record with the specified key existing in the dataset
Update:Success	count of successful record updates
Stream:Request	count of record streams obtained from the dataset
Stream:Failure	count of unsuccessful stream requests

Note:

Each of the statistics listed above is a simple counter but each has an associated rate, for instance, `Get:Success:Rate`, whose units are gets per second.

The following latencies are also available:

Dataset:GetLatency	<p>latencies of get operations. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Dataset:GetLatency#50 ■ Dataset:GetLatency#95 ■ Dataset:GetLatency#99 ■ Dataset:GetLatency#100
Dataset:AddLatency	<p>latencies of add operations. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Dataset:AddLatency#50 ■ Dataset:AddLatency#95 ■ Dataset:AddLatency#99 ■ Dataset:AddLatency#100
Dataset:UpdateLatency	<p>latencies of update operations. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Dataset:UpdateLatency#50 ■ Dataset:UpdateLatency#95 ■ Dataset:UpdateLatency#99 ■ Dataset:UpdateLatency#100
Dataset>DeleteLatency	<p>latencies of delete operations. 4 different percentiles are returned: median, 95th, 99th and maximum.</p> <ul style="list-style-type: none"> ■ Dataset>DeleteLatency#50 ■ Dataset>DeleteLatency#95 ■ Dataset>DeleteLatency#99 ■ Dataset>DeleteLatency#100

These dataset operation statistics are sent to the TMS periodically by each dataset client and then are available to view in the TMC.

There are two TMC panels that show visualizations of the dataset operation statistics:

- Overview panel: displays the latest statistics for datasets and dataset instances in a tabular layout.
- Charts panel: presents a historical view of dataset statistics over a period of time via a graphical layout.

Note:

A dataset (e.g. `dataset1`) can have multiple instances and you can distinguish each dataset instance by its name, which will end in a dash '-' followed by a number. Thus an instance of `dataset1` could be named: `dataset1-1`.

The TCStore Overview Panel

The TCStore overview panel allows you to see the latest real time statistics for all of your datasets. In addition, the datasets can be filtered by:

- Dataset name
- Clients (Dataset instances in your application)

Note:

Note that the top highlighted dataset row is an aggregation of all the dataset instance statistics.

Filter by Dataset Name

The **Datasets** dropdown option allows you to view a selected dataset. The resulting view will show all dataset instances for the selected dataset.

Screenshot: TCStore Overview Panel, Datasets dropdown.

Instance Name	Add:Failure	Add:Success	Delete:Failure	Delete:Success	Get:Failure	Get:Success
MyDataset-1 ← aggregate of all instances	506.0	20.2K	476.0	3.3K	509.0	2.3K
MyDataset-1-1 ← instance	93.0	20.0K	100.0	668.0	95.0	456.0
MyDataset-1-2 ← instance	102.0	42.0	94.0	665.0	116.0	465.0
MyDataset-1-3	115.0	43.0	94.0	650.0	95.0	469.0
MyDataset-1-4	97.0	44.0	85.0	689.0	114.0	433.0
MyDataset-1-5	99.0	50.0	103.0	671.0	89.0	466.0

Filter by a single client

The **Clients** dropdown option allows you to filter by a particular client or all clients. When selecting a single client the result will only include the dataset and dataset instances of the chosen client.

Screenshot: TCStore Overview Panel, Clients dropdown.

Instance Name	Add:Failure	Add:Success	Delete:Failure	Delete:Success
MyDataset-1	6.8K	35.9K	6.8K	28.1K
MyDataset-1-1	1.3K	23.2K	1.4K	5.7K
MyDataset-1-2	1.4K	3.2K	1.4K	5.5K
MyDataset-1-3	1.4K	3.1K	1.3K	5.7K
MyDataset-1-4	1.4K	3.2K	1.4K	5.6K
MyDataset-1-5	1.4K	3.2K	1.3K	5.6K

Filter by All Clients

When selecting all clients from the dropdown list the view will slightly change to show all clients for a particular dataset. In this layout the highlighted top row shows the aggregated statistics from all clients for the selected dataset instance. The row directly below the aggregated dataset statistics displays the client identifier, which can be expanded to show statistics for every dataset instance on that client.

Screenshot: TCStore Overview Panel, all clients of a particular dataset.

Instance Name	Add:Failure	Add:Success	Delete:Failure	Delete:Success	Get:Failure	Get:Success
MyDataset-1	540.0	20.2K	507.0	3.6K	536.0	2.4K

Instance Name	Add:Failure	Add:Success	Delete:Failure	Delete:Success	Get:Failure	Get:Success
MyDataset-1-1	0	1.0	0	12.0	0	2.0
MyDataset-1-2	1.0	1.0	0	11.0	2.0	4.0
MyDataset-1-3	0	1.0	1.0	13.0	1.0	9.0
MyDataset-1-4	3.0	0	3.0	7.0	0	5.0
MyDataset-1-5	3.0	2.0	3.0	8.0	0	0

Note:

Additional Grid Features


- Temporarily re-arrange grid columns using drag-and-drop
- Sort on columns
- Scroll horizontally if the grid columns overflow the available space

Use the Export to Excel () icon to download a spreadsheet of the currently displayed values.

Use the Filter dataset statistics () icon to select which statistics to display.

You can also filter stats by name, using space-separated terms.

Screenshot: Overlay, dataset statistics to show




The TCStore Charts Panel

The charts panel allows you to view dataset statistics over a period of time. Each statistic is represented in its own chart which shows its values over the last 5 minutes. This 5 minute window also constantly updates to ensure always seeing the most recent historical statistics.

Filter by dataset

The historical view can be filtered by a particular dataset, which is useful because a dataset could exist on more than one client.

Screenshot: TCStore Overview Panel, Historical View filtered by a particular dataset.




Filter by client

There is also an option to filter the history by a specific client or all clients.

- one client - you only see dataset statistics for the selected client, which is represented as a client identifier in the dropdown
- all clients - the dataset instance statistics are aggregated across all the clients

Screenshot: TCStore Overview Panel, Historical View filtered by clients.




Filter by instances

You can filter by dataset instances, either by a specific instance or all instances.

- one instance - this only shows statistics for the selected dataset instance
- all instances - this selection aggregates all dataset instance values for the selected dataset

Screenshot: TCStore Overview Panel, Historical View filtered by clients.



Use the slider (2 3 4) to set how many columns of charts you would like displayed.

Use the **Filter dataset statistics** () icon to select which statistics to display.

Use the **Take a snapshot of all charts** () icon to download a single PNG file containing the current values of all displayed charts.

Directly under each individual chart, use the **Export to PDF** () icon or the **Export to PNG image** () icon to download the chart in the selected format.

The TCStore Explorer Panel (Ad-hoc Query)

This panel is designed to run ad-hoc queries against a selected dataset in the cluster. It is not meant to replace programming Stream queries or the SQL query API. The panel's typical use case is: When you need to verify the existence of a record in an ad-hoc manner, you can run simple queries in this panel and check the result.


Querying clustered datasets create short-lived clients. These clients are tagged with `tmc-excluded` and thus won't appear in the TMC UI.

Access to the panel

From the Landing Page,

Select "Explorer" under "TCStore" button group,

Screenshot: Ad-hoc query from the landing page using the TCStore Explorer.




Copyright © - Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.



You will see the Explorer Panel (under the TCStore tab)

Screenshot: Explorer Panel, Search icon highlighted.



A dataset must be selected, as every query is executed against a particular dataset. If your cluster already contains datasets, the dropdown list will be populated. From this dropdown the dataset to be queried can be selected. If there is no dataset existing in your cluster, there will be a notification window indicating the error.

There are 2 query methods to choose from:

1. Query by Record Key
2. Query by Filters

Query by Record Key

If you know the record key you want, you can query by record key. Record key can be all valid Cell types except BYTES. Normally, the result should contain only one record.

Screenshot: Explorer Panel, entering a Record key.

TERRACOTTA Management Console

Ehcache **TCStore** Server Events MyCluster

Datasets: MyDataset-2 Record Key Query Filter

Record Key: key-24

Record Key ...	booleanCell (BO...	bytesCell (BYTES)	charCell (CHAR)	doubleCell (D...
key-24	false	< length: 4797 >	/	0.841499940

1 - 1 of 1 items

Query by Filters

If you don't know the record key, but you have some ideas about how certain Cells should look like, you can search by filters. The limitation is: You cannot filter by BYTES Cells.

Screenshot: Explorer Panel, search by filters.

TERRACOTTA Management Console

Ehcache **TCStore** Server Events MyCluster

Datasets: MyDataset-2 Record Key Query Filter

Query Filters: Match Type: All Any

Cell: charCell:CHAR equals e Cell Exists

Cell: doubleCell:DOUBLE lower than 0.2 Cell Exists

Record Key ...	booleanCell (BO...	bytesCell (BYTES)	charCell (CHAR)	doubleCell (DOU...	intCell (INT)	stringCell (STRI...
key-12751	false	< length: 10099 >	e	0.0853933680...	203394370	value-1730612500
key-16303	true	< length: 5620 >	e	0.1891863911...	436438387	value1438620...

You can add or remove filters. Each filter is targeted at a certain cell type (defined by cellName:cellType). Depending on the cell type, you have different query options. For example, if you select a STRING cell, then the operators can be "equals" and "starts with", which corresponds

to the portable TCStore API. Additionally, you can uncheck the "cell exists" option, and filter for all the records that MUST NOT contain the cell. This is different to not specifying any filter for the cell: If you do not specify any filter for "myStringCell", the returned records MAY OR MAY NOT contain "myStringCell". If any filter input field contains characters that cannot be converted to the target cell type, then an error notification will be displayed.

To simplify things, all the filters share the same connecting logic: You can choose "match All" or "match Any". Default is "match All".

When you think your filters look good, you can click "Run" button, and the query will be executed in the cluster against the selected dataset, in an optimized way. If you want to start over, click "Reset" button, it will clear all filters.

Grid Details

Screenshot of Grid Details

Record Key (LONG)	myBoolCell (BOOL)	myBytesCell (BYTES)	myStringCell (STRING)	myLongCell (LONG)
743	true	< length: 7 >	f5760aa	<u>N/A</u>
167	true	< length: 2 >	fd	<u>N/A</u>
436	true	< length: 3 >	f6fc95	<u>N/A</u>
2259	true	< length: 8 >	f6d9f27	<u>N/A</u>

1 - 4 of 4 items

The first column will always be Record Keys, with the key type shown in the header. Each cell header will also include the cell type. If the cell type is BYTES, the content is not rendered in String, as this would be meaningless for human consumption. Instead, the length of the Bytes content is displayed. If a record does not contain a certain cell, then you will see a specially formatted [underline]N/A. This is to visually differentiate from a String cell with content "N/A".

You can choose to sort records by column, by clicking any of the column headers. You can further filter records by records keys, by clicking on the "funnel" icon in Records Key column header.

The maximum number of records that can be returned is set to 100, a pre-determined value, intended to avoid the complete dataset to be returned to the client (and front-end UI). It is recommended to specify sufficient filters to narrow the return to only a couple of records.


Sample Cells and its configuration (advanced topic)

If you click on Sample Cells configuration button,



you will see the Sample Cell Configuration popup window,

Screenshot: Sample Cell Configuration



Sample Cells is an advanced feature. Use it only if you are sure that the returned records do not match your expected schema. Since TCStore datasets do not enforce any schema, you can have records with very different Cells (cell name and/or cell type). The management console has to send a query to sample some records/cells in order to get a general sense of the schema in order to parse and show results in the grid. The sampled cells may not represent the accurate schema. For example, you may have a cell called "myLongCell" but it's not sampled, so it will not show up in the grid. In such a situation, you have 2 options:

1. Re-sample cells until the schema matches
2. Add a custom cell with cell name "myLongCell" and cell type "LONG"

This UI window provides all features mentioned above. It features three major parts:

- You can specify a sample size. The larger the number, the more accurate, but the slower its performance. Default value is 20
- You can re-sample the cells and the results will be listed in "Sampled Cells" list
- You can add custom cells by specifying a cell name and cell type, and the current list of custom cells will be listed in "Custom Cells" list

Duplicated cells do not cause any problems, but we recommend to remove the duplicate cells in "Custom Cell" list and keep only what is needed.

The TCStore Query Statistics Panel

Screenshot: Query Statistics tab under TCStore.

Dataset	Query	Scan ...	Server...	Server...	Total ...
MyDataset-1	records().filter((intCell>?)).count()	FULL	380	10.1...	10.1...
MyDataset-1	records().filter((booleanCell==?)).c...	FULL	390	9.61...	9.64...
MyDataset-1	records().filter((charCell>?)).count()	FULL	396	8.08...	8.11...

The **Query Statistics** tab of the **TCStore** panel displays the top TCStore server-side query pipelines ordered by the total amount of server time they are costing.

Each query pipeline entry is comprised of the following fields:

Field	Description
Dataset	Name of the Dataset upon which the query was executed
Query	Shape of the server-side query pipeline
Scan Type	If the query pipeline took advantage of indexes (INDEXED) or needed to do a full scan (FULL)
Server Executions	Number of times the pipeline was presented to a server for execution
Server Time	Time taken (wall clock) by the servers to handle the query pipeline
Total Time	Time (wall clock) the query stream was held open

This information is not held by the TMS persistently, rather, it is obtained by querying each server.

These query statistics are refreshed automatically each time the tab is visited. The **Refresh** button at the top-left of the presentation is used to update the grid.

5 Using the Server Tab

The **Server** tab is where you can find information relating to the server-side of your application. There is currently a single panel, Resource Usage, described below.

Resource Usage Panel

The **Resource Usage** Panel displays information relating to your cache's or dataset's use of the Terracotta Server's configured resources, including *OffHeap* memory and Fast Restartable Store (*FRS*) data directories.

Screenshot: Resource Usage Panel

The screenshot shows the Terracotta Management Console interface. At the top, there is a blue header with the text "TERRACOTTA Management Console" and navigation icons for home and menu. Below the header is a navigation bar with tabs for "Ehcache", "TCStore", "Server" (which is selected), "Events", and "MyCluster".

The main content area is titled "Resource Usage" and includes a "Servers:" dropdown menu showing "stripe[1] --> terracotta-2-1" and an "ACTIVE" status indicator. Below this, there are "View By:" radio buttons for "Server Resource" (selected) and "Entity".

The resource usage is displayed in a tree view:

- OffHeap (714.6MB reserved of 12.0GB max capacity)
 - offheap-1 (268.4MB reserved of 4.0GB max capacity)
 - Ehcache Server Entities
 - offheap-2 (446.1MB reserved of 8.0GB max capacity)
 - TCStore Server Entities
 - MyDataset-1 (0B occupied of 87.6MB reserved, 3815 records)
 - MyDataset-2 (0B occupied of 95.9MB reserved, 3898 records)
- Data Directories
 - PLATFORM: [52.5MB]
 - /data/dataroots/platform/terracotta-2-1
 - dataroot-1: [690.6MB]
 - /data/dataroots/dataroot-1/terracotta-2-1
 - Ehcache RestartableStores
 - dataroot-2: [624.7MB]
 - /data/dataroots/dataroot-2/terracotta-2-1
 - TCStore RestartableStores
 - store/data [599.9MB occupied]

Use the **Servers** dropdown to view the resource usage of a particular Terracotta Server.

Use the **View By** radio buttons to toggle between a server resource or a caching entity-focused presentation, as shown below.

Screenshot: Resource Usage Panel, View By radio buttons

The screenshot shows the Terracotta Management Console interface. At the top, there is a blue header with 'TERRACOTTA Management Console' and navigation icons. Below the header is a navigation bar with tabs for 'Ehcache', 'TCStore', 'Server' (selected), 'Events', and 'MyCluster'. The main content area is titled 'Resource Usage' and shows a dropdown menu for 'Servers' with the value 'stripe[1] --> terracotta-2-1' and an 'ACTIVE' status indicator. Below this, there is a 'View By:' section with radio buttons for 'Server Resource' and 'Entity' (selected). The main content is a tree view of resource usage:

- OffHeap
 - Ehcache Server Entities
 - MyCacheManager (57.5MB occupied)
 - offheap-1 (268.4MB reserved of 4.0GB max capacity)
 - TCStore Server Entities
 - MyDataset-1 (117.0MB occupied of 87.6MB reserved, 3820 records)
 - MyDataset-2 (125.8MB occupied of 95.9MB reserved, 3884 records)
- Data Directories
 - MyCacheManager
 - default-frs-container/default-cachedata: [833.7MB] <dataroot-1>/ehcache/frs/default-frs-container/default-cachedata
 - MyDataset-1
 - store/data: [603.1MB] <dataroot-2>/store/data
 - MyDataset-2

OffHeap

The **Server Resource** view shows how caching entities are making use of the allocated OffHeap resources configured for the server. Any number of server entities can use a particular server resource simultaneously.

The **Entity** view shows how caching entities are using the OffHeap resources configured for the server.

A caching entity can be configured to make use of one or more server resources via the pools (both shared and dedicated) they carve out of those server resources. Caches can be configured to store their entries in the pools configured by their containing CacheManager. The amount of space occupied by a caching entity is displayed, along with the total amount of space reserved for the pool.

While caching entities are associated with offheap resource pools, dataset entities make use of server offheap resources directly. The amount of space occupied by the dataset entity, the total amount of space reserved for the offheap server resource, as well as the count of records stored are displayed. Further, the cell names of the indexes that have been defined on your dataset, the type of index, the amount of offheap space occupied, the count of records indexed, and the number of times the index has been used are all displayed.

Data Directories

Data directories are configured in the server configuration and your application points to these named disk areas for persistent storage.


For each defined data directory, the filesystem location and occupied space are displayed, as well as the names of the server entities utilizing the directory and amount of filesystem space each is occupying.

6 Using the Events Tab

The **Events** tab displays the contents of the event log, in reverse chronological order. The event log is a record of important activities occurring in your cluster. The TMS listens for and persists these events, meaning that if the TMS is not running no events will be stored.

Event Log

Screenshot: TMC Event Log



The screenshot shows the Terracotta Management Console (TMC) interface. At the top, there is a blue header with 'TERRACOTTA Management Console' and a home icon. Below the header, there is a navigation bar with tabs for 'Ehcache', 'TCStore', 'Server', and 'Events' (which is selected). To the right of the navigation bar, there is a 'MyCluster' label. Below the navigation bar, the 'Event Log' section is displayed. It includes a 'Show:' filter with checkboxes for 'TSA', 'Client', 'Ehcache', 'Entity', and 'TMS', all of which are checked. There are 'Clear' and 'Export' buttons to the right of the filter. The event log is presented as a table with the following columns: 'Alert', 'Date', 'Type', 'Name', and 'Detail'. The table contains three entries:

Alert	Date	Type	Name	Detail
	Sat Dec 08 2018 02:06:37 GMT+0100 (Central European Standard Time)	TSA	Server state changed	stripe[0]:terraccotta-1-0 moved from UNREACHABLE to ACTIVE
	Sat Dec 08 2018 02:06:04 GMT+0100 (Central European Standard Time)	TSA	Server state changed	stripe[1]:terraccotta-2-1 moved from UNREACHABLE to ACTIVE
	Sat Dec 08 2018 02:05:56 GMT+0100 (Central European Standard Time)	Client	Client left	97@10.244.5.147:Ehcache:MyCacheMana... disconnected from stripe[0]:terraccotta-1-0

The types of activities that generate events include:

- server state transitions
- clients connecting/disconnecting to TSA stripes
- server entities being created/destroyed
- resource constraint alerts
- and more

The event log is included in the set of diagnostics artifacts that can be downloaded from the cluster's connection panel on the Home Page, but in addition it can be downloaded separately using the **Export** button.

Events that are deemed to be critical are noted as alerts, displayed with an attention-focusing icon in the alert grid. Further, alert events are displayed in a temporary popup and the count of un-read alerts is shown in the cluster's connection panel on the Home Page.

Newly added events are shown in the grid with bold text, and reset to normal after navigating away.

An event is comprised of the following fields:

alert	Is this event deemed critical?
Timestamp	Time at which the TMS recorded the event
Type	Categorization of the event
Name	Sub-categorization
Detail	Description of the event

The maximum number of events that are saved in the TMS is controlled by the configuration property `tms.eventLogMaxRecords`, whose default value is 5000. The event log store in the TMS, as well as the event log grid, can be cleared using the **Clear** button.

7 Prometheus Integration

Terracotta logs metrics data in a form that can be understood by the open source data collection and filtering tool Prometheus, which is well-suited to cloud-based microservices environments. To log the data, Terracotta includes a servlet that gathers Prometheus-style metrics for put / get / save / delete operations. The Prometheus data can in turn be visualized by third party products that offer Prometheus integration (see the Prometheus web site at <https://prometheus.io/> for related information).

After you created a connection to a cluster, you can access a Prometheus compatible servlet at the following URL :

```
http(s)://[host]:[port]/actuator/prometheus
```

For example, assuming that the Terracotta Management Server (TMS) is available at `http://localhost:9480`, the servlet can be accessed at `http://localhost:9480/actuator/prometheus`.

Available metrics

For now, only resource usage metrics are available as Prometheus gauges (the same metrics as you would find under “[Resource Usage Panel](#)” on page 49).

All metrics are prefixed with `:sag_tc_server` and share the same labels or tags (as part of the plain text output, with one metric per line) :

- `alias` (can represent a server side cache resource name or an offheap resource name or a dataroot name or a dataset name)
- `connectionName`
- `entityName` (a technical attribute that represents the server-side entity name)
- `entityType` (a technical attribute that represents the server-side entity type)
- `server` (represents the server name)
- `stripe`
- `clusterTierManager` (caching resource only, matches the alias of the entity given by the user when connecting to a clustered cache, example: `terracotta://myhost:9410/anEntity`)

Here is the list of provided metrics :

Server specific resource usage metrics

- `sag_tc_server_offheap_allocated_memory_bytes` (Offheap memory allocated in bytes)
- `sag_tc_server_restartable_store_total_usage_bytes` (FRS usage in bytes)
- `sag_tc_server_dataroot_total_disk_usage_bytes` (Dataroot total disk usage in bytes)

Server-side caching specific resource usage metrics

- `sag_tc_server_caching_pool_allocated_size_bytes` (Caching pool allocated size in bytes)
- `sag_tc_server_caching_store_data_size_bytes` (Caching store data size in bytes)

Server-side store specific resource usage metrics

- `sag_tc_server_dataset_main_record_occupied_storage_bytes` (Total occupied storage by the dataset, in bytes - this is the sum of the 3 dataset_occupied metrics below)
- `sag_tc_server_dataset_occupied_primary_key_bytes`
- `sag_tc_server_dataset_occupied_persistent_support_bytes`
- `sag_tc_server_dataset_occupied_heap_bytes`
- `sag_tc_server_dataset_allocated_memory_bytes` (Total allocated storage by the dataset, in bytes - this is the sum of the 4 dataset_allocated metrics below)
- `sag_tc_server_dataset_allocated_primary_key_bytes`
- `sag_tc_server_dataset_allocated_persistent_support_bytes`
- `sag_tc_server_dataset_allocated_heap_bytes`
- `sag_tc_server_dataset_allocated_index_bytes`
- `sag_tc_server_dataset_record_count` (Dataset record count)
- `sag_tc_server_dataset_index_record_count` (Dataset indexed record count)
- `sag_tc_server_dataset_index_access_count` (Index access count)
- `sag_tc_server_dataset_index_occupied_storage_bytes` (Index occupied storage in bytes)

Examples

This section shows examples of metrics data written to the log file.

1. Dataroot example

```
sag_tc_server_dataroot_total_disk_usage_bytes{alias="PLATFORM",
connectionName="MyCluster",entityName="VoltronTmsConnection:MyCluster",
entityType="org.terracotta.management.entity.nms.client.NmsEntity",
server="terracotta-1-0",stripe="stripe[0]",} 6.5570422E7
```

Dataroot (named "PLATFORM") total disk usage in bytes; captured from the server named "terracotta-1-0", part of the stripe named "stripe[0]" (first stripe), for the connection named "MyCluster" : 60MB

2. Caching example

```
sag_tc_server_caching_pool_allocated_size_bytes{alias="dedicatedcache",
clusterTierManager="MyCacheManager",connectionName="MyCluster",
entityName="MyCacheManager$dedicatedcache",
entityType=
"org.ehcache.clustered.client.internal.store.InternalClusterTierClientEntity",
server="terracotta-1-0",stripe="stripe[0]",} 9404416.0
```

Server side cache store (created by the cache "dedicatedcache", under the clusterTierManager named "MyCacheManager") allocated size in bytes; captured from the server named "terracotta-1-0", part of the stripe named "stripe[0]" (first stripe), for the connection named "MyCluster" : 9.4MB

3. Dataset index example

```
sag_tc_server_dataset_index_access_count{
alias="booleanCell$$$Boolean$$$BTREE",connectionName="MyCluster",
entityName="MyDataset-1",
entityType="com.terracottatech.store.client.DatasetEntity",server="terracotta-1-0",
stripe="stripe[0]",} 1604.0
```

Dataset (named "MyDataset-1") index (on the cell named "booleanCell" of type Boolean; using the BTREE setting) access count captured from the server named "terracotta-1-0", part of the stripe named "stripe[0]" (first stripe), for the connection named "MyCluster" : 1604 times

Getting the metrics directly from the servlet

Using any HTTP client, such as curl :

```
curl http(s)://[host]:[port]/actuator/prometheus
```

You can craft a regular expression to search for specific metrics :

```
curl -s http(s)://[host]:[port]/actuator/prometheus |
grep "sag_tc_server.*{.*} .*"
```

Security concerns

This servlet is accessible to anybody if your TMC is not secured, but only to authenticated users if you set up authentication and authorization.

Prometheus and Grafana integration

There are several ways to let Prometheus grab the metrics available at `http(s)://[host]:[port]/actuator/prometheus`.

For example, if you deployed the TMC using Kubernetes, and you created a service for it, you can simply add this YAML configuration to your service manifest to have Prometheus reading the metrics regularly :


```

metadata:
  name: tmc
  annotations:
    prometheus.io/scrape: 'true'
    prometheus.io/path: '/actuator/prometheus'

```

Prometheus querying

When you have successfully installed and deployed Prometheus, you'll be able to choose the Terracotta cluster metrics.




You can also precisely choose which metrics you want to display using the labels and the Prometheus query language (see <https://prometheus.io/docs/prometheus/latest/querying/basics/> for related information). For example, if you only want the FRS usage for the "dataroot-2" dataroot on the server "terracotta-1-0", you can use the following Prometheus query :


```
sag_tc_server_restartable_store_total_usage_bytes{
alias=~'dataroot-2.*',server='terracotta-1-0'}
```

Grafana dashboards and panels

From Grafana, after you add a new Prometheus datasource, you can add some custom panels, using the metrics names as well as using Prometheus querying (see previous chapter).

To make things easier, we provided a Dashboard template on the Grafana website at <https://grafana.com/dashboards/10215>; just import it and start monitoring resource usage from Grafana !





8 Performance Considerations

The TMC application being both a statistics database and an application server on the same process, we strongly advise you to allocate at least 2GB of offheap memory (`tms.offheapSizeMB=2048`). Also, this memory amount will depend on how many statistics are stored and whether the persistence mode is `INMEMORY` or `HYBRID` (see below).

The amount of statistics the TMC can store depends on the number of objects in your cluster (not the number of clients). The more objects (caches or datasets) in your clients, the more memory/disk the TMC internal datastore will require.

For example, with the following configuration and a collector interval set to 30 seconds, and a cluster of 60 objects (mix of datasets and caches), the TMC's offheap memory will typically grow about 180MB per hour.

```
tms.offheapSizeMB=2048
tms.persistenceMode=INMEMORY
```

There are several parameters you can control depending on your cluster size and the quantity of history you want to keep.

- `Collector Interval`: In TMC, this controls how frequently the statistics will be collected and stored in the TMC database.
- `tms.persistenceMode`: `INMEMORY` will keep all the data in offheap memory (and disk). This speeds up the queries but requires a lot of memory if you want to keep a large history. `HYBRID` will keep only the indexes in memory. The query speed will depend on the disk speed.
- `tms.offheapSizeMB`: This controls the maximum offheap memory allowed for the TMC. `HYBRID` requires less memory than `INMEMORY`.
- `tms.offheapThreshold`: This is the threshold (as a percentage of `tms.offheapSizeMB`) above which the TMC will remove the oldest statistics from the database. This purging process is required so that the TMC does not run out of memory. 80% is a good value for 2GB of offheap memory when using `INMEMORY`.
- `tms.statisticsMaxAgeMinutes`: This parameter controls the maximum time (in minutes) the statistics will remain in the database. Statistics stored before this delay get removed. This parameter has an effect on the time frames you will be able to see in the TMC UI, assuming you have enough memory and/or disk space to store the statistics for this period of time.

9 Migrating the Terracotta Management Console

If you install a new 10.3 Terracotta Management Console over an existing installation of a previous 10.2 Terracotta Management Console , you can expect a data migration to happen.

This data migration will preserve your connections, so that you won't need to re-create them in your new 10.3 Terracotta Management Console instance.

There is no further interaction needed, except to ensure that the property

```
tms.storageFolder
```

still points to the same folder as your previous 10.2 installation did.

When starting the 10.3 Terracotta Management Console for the first time, the migration process will produce a log entry similar to the following example:

```
WARN c.t.m.s.migration.MigrationService - Backing up data before migration
WARN c.t.m.s.migration.MigrationService - Done backing up data to
/opt/softwareag/management-console/data-10.2-for-test.backup-1534351786349
WARN c.t.m.s.migration.MigrationService - Migrating data from 10.2 to 10.3
WARN c.t.m.s.migration.MigrationService - Finished migrating data from 10.2 to 10.3
```

This indicates that the migration process was successful.

