

Terracotta Server Administration Guide

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 © 2010-2020 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 <http://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 <http://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 <http://softwareag.com/licenses/> and/or in the root installation directory of the licensed product(s).

Document ID: TC-SRV-AG-107-20201015

Table of Contents

About This Documentation	5
Online Information and Support.....	6
Data Protection.....	6
1 Cluster Architecture	7
2 Active and Passive Servers	9
3 Logical Server States	13
4 Clients in a Cluster	15
5 Electing an Active Server	17
6 Failover	19
7 Starting and Stopping the Terracotta Server	21
8 Safe Cluster Shutdown and Restart Procedure	27
9 Configuration Terms and Concepts	29
10 Planning a Configuration	35
11 The Terracotta Configuration File	41
12 Config Tool	45
Introduction.....	46
Performing cluster activation and topology changes.....	48
Performing configuration changes.....	51
Diagnosing and fixing problems.....	54
13 Parameter Substitution	57
14 Configuring the Terracotta Server	59
15 System Recommendations for Hybrid Caching	63

16 System Recommendations for Fast Restart (FRS)	65
17 Failover Tuning	67
18 Connection Leasing	73
19 Cluster Tool	75
20 Licensing	89
21 Backup, Restore and Data Migration	91
Overview of Backup and Restore.....	92
Data Directory Structure.....	92
Online Backup.....	93
Offline Backup.....	94
Restore.....	95
Data Migration of Ehcache data.....	96
Technical Details.....	97
22 Migrating from older Terracotta versions to 10.7	99
23 Restarting a Stripe	101
24 IPv6 support in Terracotta	103
25 SSL / TLS Security Configuration in Terracotta	107
Security Core Concepts.....	108
Cluster Security.....	113
TMS Security.....	119
LDAP Properties.....	123
SSL / TLS Troubleshooting guide.....	126
26 Terracotta Server Migration from BigMemory to Terracotta	135
27 Using Command Central to Manage Terracotta	137
28 Terracotta in Network Environments with Subnets	141

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 Cluster Architecture

The Terracotta cluster can be viewed topologically as a collection of *clients* communicating with a Terracotta Server Array (TSA).

The server array is composed of one or more logically independent stripes. The total storage and computing capacity of the TSA can be increased with the addition of more stripes.

A stripe is composed of one or more servers. Each stripe contains a single *active server* and zero or more *passive servers*. These stripe members share a common configuration in order for any one of them to fill the "active" role. Refer to the section [“The Terracotta Configuration File” on page 41](#) for more details.

"Scale-up" is achieved by configuring the servers to utilize more storage (e.g. memory) from the machines that they are deployed on.

"Scale-out" is achieved by adding more stripes to the TSA.

Greater levels of "HA" (high availability) are achieved by adding more members and/or voters to each stripe.

For more information on clients, active servers and passive servers, see the sections [“Clients in a Cluster” on page 15](#) and [“Active and Passive Servers” on page 9](#).

TSA Topologies

There are multiple types of TSA topology, each offering different resource and availability capabilities.

TSA Topology	Description
Single-server	<p>This is a TSA which consists of one stripe containing a single server. This server is always the active server.</p> <p>This scheme offers the least amount of both resource and availability capabilities:</p> <ul style="list-style-type: none">■ If this server should become unavailable, your client end-points will fail to operate.■ The resource services exposed to your clients are limited to those of the underlying server JVM and OS.

TSA Topology	Description
High-availability	<p>This refers to a TSA where each stripe consists of at least two servers. In addition to the active server there will be at least one passive server. The stripe may continue operating in the event of an active server failure, as long as at least one passive server is available.</p> <p>Note that stripes do not share passive servers, so each stripe will need at least one passive server to possess high-availability.</p> <p>Depending upon your overall topology, the use of external voters may also help achieve greater levels of HA.</p> <p>See “Failover” on page 19 and “Failover Tuning” on page 67 for more discussion on HA topics.</p>
Multi-stripe	<p>Multi-stripe refers to a TSA that consists of multiple independent stripes.</p> <p>This scheme offers the ability for increased storage and computation resources, with each stripe contributing to the available total amount of storage.</p> <p>For a multi-stripe TSA to possess high-availability, each stripe must consist of more than one server. This setup offers the maximum of both resource and availability capabilities.</p>

Client perspective

Each client is logically independent of other clients. It sees the TSA as a collection of one or more stripes. It connects to the active server of each stripe in order to issue messages to the cluster.

Stripe perspective

Each stripe is logically independent of other stripes in the TSA. Each stripe member only concerns itself with the clients connected to it and its sibling servers.

Specifically, the active server is the key point in each stripe: each stripe has exactly one active server and it is this server which interacts directly with each connected client and each passive server within the same stripe.

2 Active and Passive Servers

Introduction

Terracotta Servers exist in two modes, *active* and *passive*. The description of each mode is given below.

Active servers

Within a given stripe of a cluster, there is always an active server. A server in a single-server stripe is always the active server. A multi-server stripe will only ever have one active server at a given point in time.

The active server is the server which clients communicate with directly. The active server relays messages on to the *passive servers* independently.

How an active server is chosen

When a stripe starts up, or a failover occurs, the online servers perform an *election* to decide which one will become the active server and lead the stripe. For more information about elections, see the section [“Electing an Active Server” on page 17](#).

How clients find the active server

Clients will attempt to connect to each server in the stripe, and only the active server will accept the connection.

The client will continue to only interact with this server until the connection is broken. It then attempts the other servers if there has been a *failover*. For more information about failover, see the section [“Failover” on page 19](#).

Responsibilities of the active server

The active server differs from passive servers in that it receives all messages from the clients. It is then responsible for sending back responses to the calling clients.

Additionally, the active server is responsible for replicating the messages that it receives on the passive servers.

When a new server joins the stripe, the active server is responsible for synchronizing its internal state with the new server, before telling it to enter a standby state. This state means that the new server is now a valid candidate to become a new active server in the case of a failover.

Passive servers

Any stripe of a cluster which has more than one running server will contain passive servers. While there is only one *active* server per stripe, there can be zero, one, or several passive servers.

Passive servers are also referred to as "mirrors", because they contain a copy of the data that is present within the active server.

Passive servers go through multiple states before being available for failover:

<i>UNINITIALIZED</i>	This passive server has just joined the stripe and has no data.
<i>SYNCHRONIZING</i>	This passive server is receiving the current state from the active server. It has some of the stripe data but not yet enough to participate in <i>failover</i> .
<i>STANDBY</i>	This passive contains the stripe data and can be a candidate to become the active, in the case of a failover.

Passive servers only communicate with the active server, not with each other, and not with any clients.

How a server becomes passive

When a stripe starts up and a server fails to win the election, it becomes a passive server.

Additionally, newly-started servers which join an existing stripe which already has an active server will become passive servers.

Responsibilities of the passive server

The passive server has far fewer responsibilities than an active server. It only receives messages from the active server, not communicating directly with other passive servers or any clients interacting with the stripe.

Its key responsibility is to be ready to take over the role of the active server in the case that the active server crashes, loses power/network, or is taken offline for maintenance/upgrade activities.

All the passive server does is apply messages which come from the active server, whether the initial state synchronization messages when the passive server first joined, or the on-going replication of new messages. This means that the state of the passive server is considered consistent with that of the active server.

Lifecycle of the passive server

When a passive server first joins a stripe and determines that its role will be passive, it is in the *UNINITIALIZED* state.

If it is a *restartable* server and also discovers existing data from a previous run, it makes a backup of that data for safety reasons. Refer to the section ["Clearing Automatic Backup Data" on page 11](#) for more details.

Refer to the section *Restarting a Stripe* in the *Terracotta Server Administration Guide* for information on the proper order in which to restart a *restartable* stripe.

From here, the active server begins sending it messages to rebuild the active server's current state on the passive server. This puts the passive server into the *SYNCHRONIZING* state.

Once the entire active state has been synchronized to the passive server, the active server tells it that synchronization is complete and the passive server now enters the *STANDBY* state. In this state, it receives messages replicated from the active server and applies them locally.

If the active server goes offline, only passive servers in the *STANDBY* state can be considered candidates to become the new active server.

Clearing Automatic Backup Data

After a passive server is restarted, for safety reasons, it may retain artifacts from previous runs. This happens when the server is restartable, even in the absence of restartable cache managers. The number of copies of backups that are retained is unlimited. Over time, and with frequent restarts, these copies may consume a substantial amount of disk space, and it may be desirable to clear up that space.

Backup rationale: If, after a full shutdown, an operator inadvertently starts the stripe members in the wrong order, this could result in data loss wherein the new active server initializes itself from the, possibly, incomplete data of a previous passive server. This situation can be mitigated by (1) ensuring all servers are running, and (2) the cluster is quiesced, prior to taking the backup. This ensures that all members of the stripe contain exactly the same data.

Clearing backup data manually: The old fast restart and platform files are backed up under the server's data directories in the format `terracotta.backup.{date&time}/ehcache/` and `backup-platform-data-{date&time}/platform-data` respectively. Simply change to the data root directory, and remove the backups.

It may be desirable to keep the latest backup copy. In that case, remove all the backup directories except the one with the latest timestamp.

3 Logical Server States

Possible server states

A server could be in any one of the following logical server states:

STARTING	server is starting
UNREACHABLE	server is unreachable
UNKNOWN	server state is unknown
UNINITIALIZED	server has started and is ready for election
SYNCHRONIZING	server is synchronizing its data with the current active server
ACTIVE	server is active and ready to accept clients
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	server is passive and ready for replication
START_SUSPENDED	server startup is suspended for all of its peers to come up
PASSIVE_SUSPENDED	server is passive but blocked in the election process (consistency mode)

4 Clients in a Cluster

Within the overall structure of the cluster, the clients represent the application end-points. They work independently but can communicate through the *active* servers of the stripes to which they are connected.

Note that a client only ever interacts with an active server, never directly communicating with a *passive* server.

In a single-stripe cluster, each client is connected to the active server of that stripe. In a multi-stripe cluster, each client is connected to the active server of *each* stripe, interacting with them quasi-independently.

Within the logical structure of the cluster, the client isn't the process making the connection, but the connection itself. This means that a single JVM opening multiple connections to the same stripe will be seen by the stripe as multiple, independent clients.

How a client finds an active server

When establishing a connection to a stripe, the client must find the active server. It does this by attempting to connect to each server in the stripe, knowing that only the active server will not reject the connection attempt.

How a client handles failover or restart

If an active server to which a client is attached goes offline, the client will attempt to reconnect to one of the other servers in the stripe, if there are any. This is similar to what happens during its initial connection.

Note that there is no default time-out on this reconnection attempt. In the case that each stripe member is unavailable, this means that it is possible for all clients to wait, blocking their progress, until a server is restarted, potentially days later.

5 Electing an Active Server

When a new stripe comes online, the first thing the servers within it need to do is elect an *active* server which will coordinate the *client* interactions and *passive* servers within the stripe.

Additionally, if the active server of an existing stripe goes offline, the remaining passive servers need to elect a replacement active server. Note that only passive servers in the *STANDBY* state are candidates for this role. For related information, see the section [“Failover” on page 19](#).

In either of these situations, the servers involved address this problem by holding an election.

A server that is started up from rest needs to get votes from all of its peer servers to get elected as an active server. In contrast, votes from a smaller set of peers are sufficient for a running *PASSIVE-STANDBY* server to become elected as an active server.

If for some reason, not all servers in a stripe can be started up, you can still forcefully get a candidate server elected as active using the cluster tool. For more information about this manual intervention using the cluster tool, see the section *The promote command* of the [“Cluster Tool” on page 75](#).

High-level process

In an election, each server will construct a "vote" which it sends to the other involved servers. The vote with the highest score can be determined statically, so each server knows it has agreement on which server won the election.

In the case of a tie, the election is re-run until consensus is achieved.

Vote construction

The vote is a list of "weights" which represent the factors which should be considered when electing the most effective active server. The list is ordered in a way that the next element is only considered if the current element is a tie. This allows the earlier elements of the vote to be based around important concepts (such as how many transactions the server has processed), followed by concrete concepts (such as server up-time), then ending in more arbitrary concepts designed to break edge-case ties (such as a randomly generated number).

6 Failover

In a high-availability stripe, the failure of a single server represents only a small disruption, but not outright failure, of the cluster and the *client* operations (for related information on high availability, see the section “[Cluster Architecture](#)” on page 7).

In the case of a failing passive server, there is no disruption at all experienced by the clients.

In the case of a failing active server, however, there is a small disruption of client progress until a new active server is elected and the client can reconnect to it. *Failover* is the name given to this scenario.

Client Reconnect Window

When a failover happens, the clients connected to the previous active server automatically switch to the new active server. However, these clients have a limited window of time called the *client reconnect window* to complete the failover (120 seconds, by default). The new active server will stop processing any client requests until all the previously known clients connect back or until this window expires. This could cause all the clients to stall even if a single client fails or takes too long to fail over to the new active server.

If clients fail to connect back to the new active server within the reconnect window, the server will consider them unreachable and will continue processing requests from the connected clients. Clients reconnecting after the reconnect window will be rejected by the server and they will rejoin the cluster as a new client by establishing a new connection.

This reconnect window can be configured in the config file or during startup using the `client-reconnect-window` property.

Server-side implications

Once all clients have reconnected (or the reconnect window closes), the server will process all re-sent messages it had seen before for which the client had not been notified of completion.

After this, message processing resumes as normal.

Client-side implications

Clients will experience a slight stall while they reconnect to the new active server. This reconnection process involves re-sending any messages the client considers to be in-flight.

After this, client operations resume as normal.

7 Starting and Stopping the Terracotta Server

Starting the Terracotta Server

The command line script to start the Terracotta Server is located in the `server/bin/` directory of the server kit. UNIX users use `start-tc-server.sh` while Windows users use `start-tc-server.bat`. All arguments are the same for both.

Options to the script

1. `-port`
Port to be used for this node. Default: 9410.
2. `-bind-address`
Bind address to be used for the node port. Default: 0.0.0.0.
3. `-group-port`
Port to be used for intra-stripe communication. Default: 9430.
4. `-group-bind-address`
Bind address to be used for node group port. Default: 0.0.0.0.
5. `-hostname`
Host name for this node. Must be a valid DNS name, or an IP address. Default: %h.
6. `-name`
Name to be used for this node. Needs to be unique in the cluster. Default: a randomly generated string.
7. `-public-hostname`
Public host name for this node. Needs to be set on all the nodes in the cluster. See [“Terracotta in Network Environments with Subnets” on page 141](#) for more details about this feature.
8. `-public-port`
Public node port for this node. Needs to be set on all the nodes in the cluster. See [“Terracotta in Network Environments with Subnets” on page 141](#) for more details about this feature.

9. `-config-dir`

Directory containing configuration bookkeeping information generated by the node. Default: `%H/terracotta/config`.

10. `-backup-dir`

Directory to be used to contain cluster backups. Needs to be set on all nodes in the cluster. See the section [“The “backup” Command” on page 79](#) of the cluster tool for more details about performing backups.

11. `-log-dir`

Directory to be used to contain logs for this node. Default: `%H/terracotta/logs`.

12. `-metadata-dir`

Directory to be used to contain server persistence data. Default: `%H/terracotta/metadata`. See the section [“Server Persistence” on page 60](#) for more details about this setting.

13. `-repair-mode`

Whether to start the node in repair mode. Default: `false`.

14. `-client-lease-duration`

Lease duration for the client connections. Default: 150s. Needs to be the same throughout the cluster. See [“Connection Leasing” on page 73](#) for more details about this property.

15. `-client-reconnect-window`

Time window for clients to reconnect to a new Active server after failover. Default: 120s. Needs to be the same throughout the cluster.

16. `-cluster-name`

Name to assign to the cluster this node will become a part of. Needs to be the same throughout the cluster.

17. `-config-file`

Config file to start the node with. See the section [“The Terracotta Configuration File” on page 41](#) for more details about this property.

18. `-failover-priority`

The failover priority setting to be used. Valid values are `availability` and `consistency`. Needs to be the same throughout the cluster. This is the only mandatory option if the node is started using console parameters. See [“Failover Tuning” on page 67](#) for more details about this property.

19. `-data-dirs`

Directory to contain client data. Default: `main:%H/terracotta/user-data/main`. Data directory names needs to be the same throughout the cluster - the disk locations could vary. See [“Configuring the Terracotta Server” on page 59](#) for more details about this property.

20. `-offheap-resources`

Offheap resources to be used. Default: `main:512MB`. Needs to be the same throughout the cluster. See [“Configuring the Terracotta Server” on page 59](#) for more details about this property. Offheap resources is the only property which is validated against the license. The total `offheap-resources` for the cluster (i.e. `offheap-resources` for all stripes summed up) should be within the license limit. See [“Licensing” on page 89](#) for more details about licensing.

21. `-audit-log-dir`

Directory containing the node's security audit logs. Needs to be set on all the nodes. Needs some form of security to be enabled.

22. `-authc`

Security authentication setting to be used. Valid values are `file`, `ldap` and `certificate`. Needs to be the same throughout the cluster.

23. `-security-dir`

The security root directory for this node.

24. `-ssl-tls`

Whether to enable SSL/TLS based security. Default: `false`. Needs to be the same throughout the cluster.

25. `-whitelist`

Whether to enable IP whitelist security. Default: `false`. Needs to be the same throughout the cluster.

See [“Cluster Security” on page 113](#) for more details about configuring security.

Use cases

The server startup script can be used in different ways, depending on the needs and convenience of the server admin:

Starting nodes with console parameters

A server can be started in unconfigured node, with console parameters alone (i.e. without the need of having any config files).

Example:

```
./start-tc-server.sh -config-dir /data/tc/node1
  -port 9410 -group-port 9430 -name node1 -failover-priority consistency
  -offheap-resources primary:250GB,caching:100GB
```

Note: the `config-dir` in question must be empty, as it only makes sense to use the command with these parameters to load a new server with its initial set of configuration properties.

After startup, it will then be necessary to execute the `config tool attach` and `activate` commands respectively to define the cluster topology and to make the cluster ready for client operations.

See the sections [“Adding nodes using the attach command” on page 49](#) and [“Activating a cluster using activate command” on page 48](#) for related details.

Before activating, the config tool could also be used to set various other configuration options on the node.

Starting nodes with config file

An unconfigured server can be also be started with a parameter specifying a config file which was exported from an existing cluster, or constructed by hand. See the section [“Exporting cluster configuration” on page 53](#) for related information.

Example:

```
./start-tc-server.sh -config-dir /data/tc/node1  
-config-file myCluster.properties
```

In this case it is not necessary to execute the config tool attach command since the cluster topology is already defined in the config file. The execution of the activate command, however, would still be required to make the cluster ready for client operations.

Note:

The config-dir in question must be empty, as it only makes sense to use the command with these parameters to load a new server with its initial set of configuration properties.

Note:

it is important to understand that the configuration file used in this command is not what will be used by the server each time it starts. Rather, it is simply a convenience for passing the initial configuration properties to the server, instead of specifying each config property on the command line and/or setting it via config tool. Once the config file is read, the server's internal configuration is stored in its config-dir. The config-dir contains the configuration that will be used on subsequent restarts.

Starting nodes with config dir

A previously configured server can be started with the specification of a configuration directory (config-dir), in activated mode directly.

Example:

```
./start-tc-server.sh -config-dir /data/tc/node1
```

This is used when the server was restarted after having been activated previously, but can also be used on a fresh server startup after completing the migration process from an older Terracotta cluster. See the section [“Migrating from older Terracotta versions to 10.7” on page 99](#) for related details

Note that when you start a server this way (which will be the most common way, over time), any other parameters passed will be effectively ignored, because the server will use its internal configuration.

After starting a node this way, it is not necessary to run the config tool attach or activate commands since the cluster topology and configuration are ready.

Environment variables read by the script

- `JAVA_HOME` - Points to the JRE installation which the server should use (the Java launcher in this JRE will be used to start the server).
- `JAVA_OPTS` - Any additional options which should be passed to the underlying JVM can be passed via this environment variable and they will be added to the Java command line.

Stopping the Terracotta Server

If your server is not running in a Terracotta cluster, you can use the standard procedure offered by your operating system to terminate the server process.

If you are running the server as part of a Terracotta cluster, you can safely shut down all servers in the cluster using the cluster tool. See the section [“The "shutdown" Command” on page 81](#) of the cluster tool for details.

8 Safe Cluster Shutdown and Restart Procedure

Although the Terracotta Server Array is designed to be crash tolerant, like any distributed system with HA capabilities, it is important to consider the implications of shutting down and restarting servers, what sequence that is done in, and what effects that has on client applications and potential loss of some data.

The safest shutdown procedure

For the safest shutdown procedure, follow these steps:

1. Shut down all clients and ensure no critical operations such as backup are running on the cluster. The Terracotta client will shut down when you shut down your application.
2. Use the `shutdown` command of the cluster tool to shut down the cluster.

If you want to partially shut down a stripe with passive servers configured, you can use the partial shutdown commands provided by the cluster tool. See the section [“Cluster Tool” on page 75](#) for details.

The safest restart procedure

To restart a stripe for which the failover priority is *consistency*, servers can be started up in any order as it is guaranteed that the last active server is re-elected as the active server, thus preventing data loss. This is guaranteed even if there are multiple former active servers in the stripe at the time of shutdown (for example, one active server and one or more suspended active servers or former active servers that were shut down, decommissioned or had crashed).

However, if the failover priority is *availability*, restarting the servers in any random order might result in data loss. For example, if an older active server is started up before the last active server, it could win the election and become the active server with its old data. To avoid such data loss scenarios, the last known active server must be restarted first. All other servers must be started up after this last known active server becomes the active server again.

However, if you do not know the most recent active server at the time of restart and still want to restart the stripe safely without data loss, it can still be done by starting all the servers in that stripe using the `--consistency-on-start` option of the server startup script. When the servers are started up using this option, they will wait for all peer servers to come up and then elect the most recent active server as the new active server.

If there are multiple active servers at the time of shutdown, which can happen if the failover priority of the cluster is *availability*, one of them will be chosen automatically on restart. This choice is made based on factors like the number of clients connected to those servers at the time of shutdown, the server that was started up first, etc.

Considerations and implications of not following the above procedure

Facts to understand:

- Servers that are in "active" status have the "master" or "source of full truth" copy of data for the stripe they belong to. They also have state information about in-progress client transactions, and client-held locks.
- Mirror servers (in "passive standby" state) have a "nearly" up to date copy of the data and state information. (Any information that they don't have is redundant between the active server and the client.)
- If the active server fails (or is shut down purposely), not only does the standby server need to reach active state, but the clients also need to reconnect to it and complete their open transactions, or data may be lost.
- A Terracotta Server Array, or "Cluster" instance has an identity, and the stripes within the TSA have a "stripe ID". In order to protect data integrity, running clients ensure that they only "fail over" to servers with matching IDs to the ones they were last connected to. If cluster or stripe is completely "wiped" of data (by purposely clearing persisted data, or having persistence disabled and having all stripe members stopped at the same time), that will reset the stripe ID.

What happens if clients are not shut down

If clients are not shut down:

- Client applications will continue sending transactions (data writes) to the active server(s) as normal, right up until the active server is stopped. This may leave some successful transactions unacknowledged, or falsely reported as failed to the client, possibly resulting in some data loss.
- Clients will continue to try and connect and when the server is restarted, the clients will fail the current operation and enter a reconnect path to try and complete the operation. When clients enter a reconnect path, it is left to the client to ensure idempotency of the ongoing operation as the operation might either have been made durable just before shutdown or it may have been missed during shutdown.

What happens if the active server is shut down explicitly

If the active server is shut down first:

- Before shutting down any other servers, or restarting the server, ensure that you wait until any other servers in the stripe (that were in 'standby' status) have reached *active* state, and that any running clients have reconnected and re-sent their partially completed transactions. Otherwise there may be some data loss.

9 Configuration Terms and Concepts

To have a solid grasp of how to configure a Terracotta Server Array (TSA), one must first have a strong understanding of the basic concepts of what a TSA is, what it uses as resources, and how its configuration system works.

Review of Terracotta Server Array Concepts

As a quick review of high-level TSA concepts:

- A Terracotta Server Array (TSA) is composed of one or more "stripes".
- A stripe is composed of one or more Terracotta Servers
- Each stripe contributes to the total storage and computing capacity of the TSA. If there are five stripes, then each one will contain roughly one-fifth of the stored data.
- Within a stripe, one server is "active" (serves workload from clients), and any others act as "mirrors" for HA purposes.
- Because any member of the stripe may be elected as the "active" server, the configuration and system resources of all stripe members must be equivalent.

Stripes have names (which can be assigned during configuration time), and nodes (servers) that are members of the stripe also have names. These names are useful when "targeting" configuration or operational commands.

For more information on the above concepts, please review the sections [“Cluster Architecture” on page 7](#) and [“Active and Passive Servers” on page 9](#).

Server Resource Concepts

Terracotta Servers utilize resources in order to provide services and features such as network connectivity, data storage, durability, backups, etc.

Notable items that need to be configured (or considered whether the default value is appropriate) include:

- Network ports - This includes a "port" for receiving client requests, and a "group-port" for communicating with other stripe members (servers). The default values are 9410 and 9430 respectively.

- **Server metadata directory** - This is a directory where the server stores important metadata about its internal state. The default location is `<user-home-dir>/terracotta/metadata`.
- **Configuration directory** - This is a directory where the server stores its internal configuration. The default location is `<user-home-dir>/terracotta/config`.
- **Offheap storage resources** - Servers need one or more offheap (memory) resources defined in order to have space in which to store data (via Caches or Datasets). For proper operation, all servers in a cluster (TSA) need to have the same set of offheap resources defined (because Cache and Dataset configurations will reference them for use). The default is to create one offheap resource named `main` with size 512MB.
- **Data directories** - Optional, but commonly used, data directories are used for durable (persistent) storage of data. For proper operation, all servers in a cluster (TSA) need to have the same set of data directories defined (because Cache and Dataset configurations will reference them for use). The default is to create one data directory named `main` with location `<user-home-dir>/terracotta/user-data/main`.
- **Backup directory** - Used as the destination for backups.
- **Logging directory** - Used as the destination for server logs. Default location is `<user-home-dir>/terracotta/logs`.
- **Failover priority** - a cluster-wide setting that affects HA behavior when nodes are shut down or fail. A choice must be made as to whether the cluster should favor `availability` of service or `consistency` of data when situations occur that could lead to split-brain scenarios (e.g. when servers are still running but cannot communicate with each other).

For more information on about these items, please review the sections [“Config Tool” on page 45](#), [“The Terracotta Configuration File” on page 41](#), and [“Configuring the Terracotta Server” on page 59](#).

Configuration Concepts

Perhaps the most important thing to understand about a Terracotta server's configuration, is that it is stored and updated "internally" to the server, not in a human-editable file that is read each time the servers starts.

After a node has been configured and is running, everything that is needed to restart it, and get it running again (with the same configuration and internal state) is stored within the node's `config-dir`.

The mechanisms for adding to or changing the internally stored configuration of servers is therefore the focus of what needs to be understood next.

Fundamental, Required Settings

There are a few very fundamental items related to a server instance (node) that are necessary for its existence. This includes: network ports, `config-dir` and `metadata-dir`. The ports are of course used to make the node accessible. The `config-dir` is where the server will store (and later find) its internal configuration. And the `metadata-dir` is where the server will store (and later find) its internal state information.

Necessarily Equivalent Settings

Some configuration settings need to be consistent or equivalent across all nodes in a stripe and/or across all nodes in all stripes of a cluster. The reasoning for this is fairly clear and logical, if we give a few examples.

Settings that need not be (or, in some cases must not be) the same across nodes include things such as the node name. Clearly, the node name must be unique, or it wouldn't be useful for identification of the node. The network port can be any legal and available port number, and there is no need for any two servers to use the same port, though it probably is most clear if they all use the same port. (Note that they obviously must not use the same port if they are on the same host, but for other reasons (resources, and HA) it is strongly recommended not to run servers on the same host).

Some examples of settings that need to be equivalent across nodes are offheap resources and data directories. This is because these are referenced (by name) in configurations for Datasets and Caches, and are expected to exist on all servers. For example, if a user configures a new Dataset to utilize (store its data in) an offheap resource named 'primary', then, as the Dataset is created on each server member of the cluster, an offheap resource named 'primary' must be found on each, or the creation of the dataset will fail. It also makes sense that the offheap resource named 'primary' should have an identical size on each server, such that mirrors can hold a copy of all the same data the active server has, etc. For data-dirs, it is similar: when a Cache or Dataset configuration instructs the usage of a data-dir, one with that name must exist on each server in the TSA. (However, in the case of a data-dir, while the data-dir name needs to be known by all servers, the file path that it refers to does not have to be *identical* on all servers. Hence, we say that all server nodes should have the *equivalent* set of data-dirs.)

Initial Configuration Steps

The typical steps for initial configuration of a TSA are:

1. Start each server node (as unconfigured servers, they will enter 'diagnostic mode', and await configuration)
2. Provide each server with configuration settings
3. Attach nodes to each other to form stripes
4. Attach stripes to each other to form a cluster
5. Activate the cluster

The first two steps can be accomplished in one command-line, if the user specifies configuration settings as parameters to the `start-tc-server` script.

All of the steps can be accomplished in one command-line, if the user specifies a config file containing all of the settings for all nodes of the cluster. Note that such a file is only used to initialize the set of configuration properties in each node's internal configuration (stored in its config-dir) - it is never read or utilized again.

After these steps, the server nodes will restart themselves (to leave diagnostic mode) and form the configured cluster. As the servers restart, they will use the configuration stored in their

config-dir. With any subsequent restarts of a node, the user must specify (to the start-tc-server script) the location of the config-dir (or the script will attempt to use the default location).

Once the cluster is activated, some configuration properties (such as node name and config-dir) cannot be changed. Others can be changed or added later, as necessary (such as offheap resources).

Understanding The Configuration Directory and Config Tool

Configuration Directory

As previously noted, a server node's config-dir is where its internal storage, or "source of truth", for configuration is kept. The files under this directory are *not to be edited by the user*, as (for reasons that will be made more clear below) they are solely managed by the server node itself.

If you are restarting a server process, and want it to be the same server node that it was before, you need to ensure that the start-tc-server script specifies (or defaults to) the appropriate config-dir.

Config Tool

The config tool (see [“Config Tool” on page 45](#)) is used to add or modify configuration settings for servers, both before they are activated as part of a cluster, and afterward. It can also be used to see what a server's current configuration settings are, or export them for use as a template or backup for recreating clusters.

The config tool connects to server nodes and issues commands, such as to set a configuration property. Some config tool commands may target a single node, while others may target all nodes of a stripe, or all nodes of a cluster. In all cases, the server responds to the config tool's requests by reading and/or updating the configuration state files contained in the server's config-dir.

Configuration Operations and Outcomes

As noted in the previous paragraphs, a server node responds or reacts to config-tool requests by reading or updating the contents of its internal configuration (which is contained in the config-dir, once the cluster has been activated).

Because many configuration settings must be the same on all member nodes of a stripe and some must be the same on all members of the cluster, changes to configuration must be coordinated, and therefore complex outcomes are possible.

For example, if the number of "voters" for a stripe is to be changed, it is only safe for that change to go into effect if it does so at the same time on all nodes in the stripe. Otherwise, "bad things" could happen when a failover situation occurs (e.g. one node may make a bad decision as to whether or not it should move to "active" state.) Similarly the adjustment of offheap resource sizes, or the attachment of an additional node to a stripe needs to be synchronized across servers.

This coordination, or synchronization is accomplished via a two-phase commit protocol, wherein configuration changes are staged and validated on each server (to determine that the change can be successful on all servers), and then committed (or activated) on each server in the second phase.

Typically, (in most cases), the config tool and the server's internal configuration manager handle the complexities of the coordination just fine. However, if a failure occurs during a configuration

change, or one of the nodes is not running when a configuration change is made, there are possible outcomes that may require follow-up on the user's part.

Most of the non-typical cases are automatically corrected by the servers, such as when a node restarts (if it was down during the configuration change, or if it crashed after the configuration change was staged but before it was committed or rolled back). When servers that are restarting connect to the other members of their stripe, they discover whether their configuration state is out of sync, and if so, then they receive the appropriate updates from the other server(s).

Very rarely the config tool's `repair` command may need to be used to force a commit or rollback of a config change after careful inspection of the configuration state via the config tool's `get` or `export` commands.

In all cases the config tool will inform you about the success or failure of a configuration operation, and hint to you any next steps that may be necessary.

10 Planning a Configuration

To be successful, most deployments require at least a little planning before beginning the configuration process in order to avoid missteps or even the need for starting over.

If you have not already done so, please familiarize yourself with the material presented in [“Configuration Terms and Concepts” on page 29](#).

Naming and Addressing

Naming Servers and Stripes

Because Terracotta deployments typically involve at least two servers, and often very many more, you should put a little planning into how you name them. Doing so will help keep the nodes and stripes easily identifiable within configuration and management commands and monitoring views. It is not actually necessary to name them, in which case they will be assigned auto-generated names, but using names that are meaningful to you will likely be helpful.

Some things to consider when deciding upon the scheme for naming stripes and nodes:

- You may want to include within the name of a stripe or node something that hints at its purpose, such as whether it is part of a development, test, or production environment. For example "DevStripe-A".
- You may want to include within a node's name something related to the name of the host upon which the server runs. On the other hand, in dynamic/container environments you may want to purposely avoid this.
- If you expect that you'll be changing your TSA topology in the future (i.e. adding or removing stripes from the cluster, or adding or removing servers from stripes), you may want to purposely avoid using sequential numbering in the names, such as "server-1" or "stripe-1", because over time you may end up with gaps or other oddities in the numbering. This will actually work fine, but may be confusing to users who try to assemble within their minds a mental map of the topology.

As you form your cluster, the cluster itself can also be named. It makes good sense to use a name that clearly identifies its purpose, e.g. "MyApp-PROD-TSA" or "MyApp-DEV-TSA", etc.

Addressing Servers

As you plan the set of servers that you will need, and where they are to be deployed, it would be wise to make, and keep handy during the configuration process, a clear listing of which host names (or addresses) and ports will be used for each server.

You have the choice of addressing servers by hostname or by IP address. Using host names (that can be resolved by DNS) is favorable. Note that you can also specify `bind-address` (for `port`) and `group-bind-address` (for `group-port`) in case of any ambiguity of which IP address the ports will be opened on (with the default being to open the ports on all of the host's addresses).

Data Consistency and Availability

One of the most important actions in planning your configuration is that of determining which guarantees you would like to favor in the case of a fail-over situation.

Please refer to [“Failover Tuning” on page 67](#) for a full discussion of that feature.

The well-known CAP Theorem causes the need for a choice in which guarantees the TSA should sacrifice in order to preserve the others, in the case of a fail-over situation.

If you plan to store data in the TSA and have its integrity protected with priority, you should strongly consider using the consistency setting for the cluster's `failover-priority` setting.

If you plan only to cache data in the TSA, you may prefer to use the `availability` setting for `failover-priority`.

In either case, the likelihood of the cluster ever needing to resort to compromising on either data availability or data consistency can be greatly reduced by careful choices and resourcing related to High Availability.

High Availability

High Availability (HA) of the Terracotta Server Array is achieved through the use of mirror servers within each stripe, and the optional use of "voters".

When an active server is shut down or fails, other stripe members become eligible for becoming the new active server for that stripe's set of data. If there are no other members of the stripe running, then the stripe's data is not available, and that typically results in the complete unavailability of the Terracotta cluster, until the stripe is back online.

As you plan your TSA configuration, you should consider what levels of service are required, and plan the proper number of servers per stripe and any requisite external voters (to assist with tie-breaking during elections when quorum is not otherwise present).

For more information on these topics, see [“Active and Passive Servers” on page 9](#), [“Electing an Active Server” on page 17](#), and [“Failover Tuning” on page 67](#) (including discussion of External Voters).

Storage and Persistence Resources

In-Memory Storage

As part of planning for your configuration, you need to put some thought into how you will organize the storage of your data.

Typically, data is stored within "offheap resources" which represent pools of memory reserved from the underlying operating system. You configure one or more offheap resources, giving each a name and a size (such as 700MB or 512GB, etc.). After your cluster is up and running, you can create Caches and Datasets for storing your data, and as you do so, you will need to indicate which offheap resource will be used by each.

There is nothing inherently wrong with simply defining only one offheap resource and having all Datasets and Caches use it. However, some users may find it useful to be sure particular amounts of memory are reserved for particular Datasets or Caches.

Note that the total amount of memory for a particular offheap resource is actually the configured size of the resource multiplied by the number of stripes in the TSA, because the configured amount is for a given server. Thus if you configure an offheap resource name 'primary' with a size of 50GB, and your TSA has 3 stripes, then you will be able to store a total of 150GB of data (including any related secondary indexes) within the 'primary' offheap resource.

See also the topic *Necessarily Equivalent Settings* in the section "[Configuration Terms and Concepts](#)" on page 29.

Disk Storage and Persistence

Most users desire to have at least some sets of their data persisted, or in other words, durable between restarts of the servers in the TSA. Terracotta's FRS and Hybrid features provide such capability.

FRS (Fast Restartable Store) is a transaction log structured in such a way that it can be very efficiently replayed upon server restart, in order to recover all of the stored data as it existed when the server went down. (Note that passive/mirror servers would instead sync the latest state of the data from the active server). When FRS is enabled, data writes (additions, updates, deletions) are recorded in FRS, but all data reads (gets and queries) occur within memory.

Hybrid storage mode utilized FRS capabilities, but also expands storage capacity to include the disk, not just memory. In Hybrid mode, memory (offheap resources) is used to store keys, pointers/references and search indexes (for extremely fast resolution of lookups and queries), but values are read from disk, such that memory does not need to have the capacity to contain them all. Like FRS, all data is recovered in its last state when the server restarts.

For your configuration planning, you should note that both FRS and Hybrid features require a location on disk where they can store the data. Because data is written to disk when modifications occur, the speed of the disk is a major factor on the latency and throughput of Dataset and Cache operations, and the speed of server restarts. Many users find it beneficial to dedicate a highly performant file system for FRS/Hybrid data, while having the server use a different file system for storing configuration, logs, etc. Some users find it useful to have multiple file system paths (e.g. mount points) for storing different sets of data (different Caches or Datasets) both for performance and organizational (e.g. for backups) purposes.

Locations for user data storage are specified with the `data-dirs` configuration property, which can contain a comma-separated list of one or more data directories. Each `data-dir` has an identifying name, that is used in the configuration of Caches and Datasets to enable persistence of the data

that is put into them. Recall that an equivalent set of data-dirs (with the same names), should exist on all nodes of the cluster.

Your planning should consider what filesystem path(s) you will use for data persistence (if any), and what names you would like to identify each of those locations with.

See also the topic *Necessarily Equivalent Settings* in the section [“Configuration Terms and Concepts” on page 29](#).

Config and Metadata Directories

Terracotta servers require locations for storing their internal configuration (set with the `config-dir` property), and their state metadata (set with the `metadata-dir` property).

For each server instance you'll want to make sure that the locations of these are always available to the server (perhaps ideally on a local disk).

You may want to plan to name the directories after the server node's name, or similar - in order to help keep things organized and clear for yourself and others who administer the system.

Backup Directory

In order to use the data backup feature of Terracotta, you will need to configure a location for the backup to be written to. This is done with the `backup-dir` config property.

The location should ideally be performant (such that the backup files can be written quickly, with minimal impact to the server), and large enough to contain the backup to be made, plus any other backups that you may have previously made and not removed.

See also: [“Backup, Restore and Data Migration” on page 91](#).

Logging Directory

You should also put some planning into where your server's logs will be written. This is configured with the `log-dir` property.

Like the server's metadata directory and config directory, the log directory should be available to the server at all times, and you likely want to ensure that its path and name make clear sense to you and others who will be administering the system, as to which server's logs the directory contains.

Security

Your configuration planning should also consider whether you wish to enable security features on your cluster. Security features include encryption of network communications via TLS/SSL, and authentication, authorization and auditing (AAA) features.

If so, you will need to become familiar with these feature to properly plan your configuration. See also: [“Security Core Concepts” on page 108](#) and [“Cluster Security” on page 113](#).

Public Addresses

Will the clients need to address the servers differently than the servers address each other (such as due to being within a managed container environment that has an "internal" network)?

If so, you may want to review whether hostnames will resolve to legal addresses both inside and outside of the containers, and whether you need to use the `public-address` configuration setting on your servers.

See also: [“Terracotta in Network Environments with Subnets” on page 141](#).

11 The Terracotta Configuration File

This document describes the elements of the Terracotta configuration file, which is which is a Java properties file.

The config file serves to describe the initial configuration for all members of a Terracotta cluster. Refer to the section [“Cluster Architecture” on page 7](#) for details on various different TSA topologies.

A sample config file is provided in the kit under `server/conf`, which can be used as the starting point. Some samples have inline comments describing the configuration elements. Be sure to start with a clean file for your configuration.

Use cases

Unlike previous Terracotta releases, a config file isn't needed to configure or start the servers. However, a config file is handy for the following purposes:

1. Exporting and viewing the configuration of a given cluster. See the section [“Exporting cluster configuration” on page 53](#).
2. Backing up and version-controlling the configuration of a given cluster
3. Using the current configuration as a foundation to build up a new configuration
4. Importing the cluster configuration on running nodes of an unconfigured cluster. See the section [“Importing cluster configuration” on page 54](#).
5. Starting nodes using the exported config in a new cluster. See the section [“Starting and Stopping the Terracotta Server” on page 21](#).

It is important to understand that in these cases the configuration file is only used as a convenience to "feed" all of the configuration settings into unconfigured servers. Once that is done, the servers will never again utilize or reference the config file, but will instead utilize their internal configuration (stored in the server's `config-dir`).

Examples

This section describes the various properties supported in a config file.

Minimal configuration

A config file can get pretty large, especially when the cluster contains a large number of nodes. However, most configuration properties have default values, because of which the config file can be reduced in size if the default values are acceptable. The only two mandatory properties for a single-node cluster are:

```
failover-priority=availability
stripe.1.node.1.hostname=localhost
```

which specifies the `failover-priority` of the cluster to `availability`, and the `hostname` to `localhost`. The following default values are used:

Property	Default value	Comments
<code>offheap-resources</code>	<code>main:512MB</code>	Defines one offheap resource with name <code>main</code> and size <code>512MB</code> . See “Offheap Resources” on page 59 .
<code>client-lease-duration</code>	<code>150 seconds</code>	Defines the lease duration for the client connections as <code>150 seconds</code> . See “Connection Leasing” on page 73 .
<code>client-reconnect-window</code>	<code>120 seconds</code>	Defines the client reconnect time window as <code>120 seconds</code>
<code>stripe.1.node.1.name</code>	<code><randomly-generated></code>	A randomly-generated name for the node.
<code>stripe.1.node.1.hostname</code>	<code>%h</code>	Sets the host name of the node to the fully-qualified host name of the machine.
<code>stripe.1.node.1.port</code>	<code>9410</code>	Sets the port for this server process to <code>9410</code>
<code>stripe.1.node.1.bind-address</code>	<code>0.0.0.0</code>	Sets the bind address for the port to the wildcard address <code>0.0.0.0</code>
<code>stripe.1.node.1.group-port</code>	<code>9430</code>	Sets the intra-stripe communication port for this server process to <code>9430</code>
<code>stripe.1.node.1.group-bind-address</code>	<code>0.0.0.0</code>	Sets the bind address for the group port to the wildcard address <code>0.0.0.0</code>
<code>stripe.1.node.1.metadata-dir</code>	<code>%H/terracotta/metadata</code>	Sets the server persistence directory to <code>%H/terracotta/metadata</code> . See “Server Persistence” on page 60 .

Property	Default value	Comments
stripe.1.node.1.data-dirs	main:%H/terracotta/user-data/main	Defines a user data directory with name <code>main</code> and path <code>%H/terracotta/user-data/main</code> . See “Data Directories” on page 60 .
stripe.1.node.1.log-dir	%H/terracotta/logs	Sets the server logging directory to <code>%H/terracotta/logs</code>

Note:

`%h` and `%H` in the above default values point to the hostname of the machine and home directory of the current user respectively. See the section [“Parameter Substitution” on page 57](#) for more information.

Several other configuration properties are omitted (i.e. not assumed to have defaults), which are:

Property	Comments
cluster-name	The name of the cluster
whitelist	Whether to enable IP whitelist security
ssl-tls	Whether to enable SSL/TLS based security
authc	Security authentication setting to be used
security-dir	The security root directory for this node
audit-log-dir	Directory containing the node's security audit logs
backup-dir	Directory to be used to contain the backup of this node. See the section “Backup, Restore and Data Migration” on page 91 .
public-hostname	Public hostname for this node. See the section “Terracotta in Network Environments with Subnets” on page 141 .
public-port	Public port for this node. See the section “Terracotta in Network Environments with Subnets” on page 141 .

Security configuration

The following snippet demonstrates how to enable IP whitelisting and SSL/TLS based security (see the section [“SSL / TLS Security Configuration in Terracotta” on page 107](#)) along with security event auditing on a single node cluster:

```
failover-priority=availability
whitelist=true
```

```
ssl-tls=true
authc=certificate
stripe.1.node.1.hostname=localhost
stripe.1.node.1.audit-log-dir=/path/to/audit/dir
stripe.1.node.1.security-dir=/path/to/security/dir
```

High availability configuration

High-availability can be enabled by configuring more than one node in a stripe. The following snippet defines two nodes in a cluster containing a single stripe:

```
failover-priority=availability
stripe.1.node.1.hostname=localhost
stripe.1.node.1.port=9410
stripe.1.node.1.group-port=9430
stripe.1.node.2.hostname=localhost
stripe.1.node.2.port=9510
stripe.1.node.2.group-port=9530
```

Multistripe with HA configuration

Stripes can be added to a Terracotta cluster to scale it out. Additionally, high-availability in a stripe can be enabled by configuring more than one node. The following snippet defines a cluster with two stripes with two nodes each:

```
failover-priority=availability
stripe.1.node.1.hostname=localhost
stripe.1.node.1.port=9410
stripe.1.node.1.group-port=9430
stripe.1.node.2.hostname=localhost
stripe.1.node.2.port=9510
stripe.1.node.2.group-port=9530
stripe.2.node.1.hostname=localhost
stripe.2.node.1.port=9610
stripe.2.node.1.group-port=9630
stripe.2.node.2.hostname=localhost
stripe.2.node.2.port=9710
stripe.2.node.2.group-port=9730
```

12 Config Tool

■ Introduction	46
■ Performing cluster activation and topology changes	48
■ Performing configuration changes	51
■ Diagnosing and fixing problems	54

Introduction

Config tool is a command-line utility that allows administrators of the Terracotta Server Array to perform a variety of cluster management tasks. For example, the config tool can be used to:

- Activate a cluster
- Update license on a cluster
- Make configuration changes on the nodes of a cluster
- Make topology changes in the cluster

The config tool script is located in `tools/bin` under the product installation directory as `config-tool.bat` for Windows platforms, and as `config-tool.sh` for Unix/Linux.

The config tool has several options and commands, which we will detail below.

Please see the section [“Configuration Terms and Concepts” on page 29](#), which will give you a valuable overview for better understanding what the config tool commands do, and why and when you would use them.

List of settings

Legend:

- **Visibility:** whether the setting change is visible at `RUNTIME` or after a `RESTART`. `READ ONLY` means that the setting, once defined, cannot be changed.
- **Applicability:** whether a setting can be changed on a particular `NODE`, on all nodes of a `STRIPE`, or on all the nodes of a `CLUSTER`.
- **Requirement:** whether a setting can be changed if `ALL NODES` are up or only if at least all `ACTIVES` (1 active per stripe) are up.

Setting Name	Visibility	Applicability	Requirement
name	READ ONLY	N/A	N/A
hostname	READ ONLY	N/A	N/A
port	READ ONLY	N/A	N/A
group-port	READ ONLY	N/A	N/A
bind-address	READ ONLY	N/A	N/A
group-bind-address	READ ONLY	N/A	N/A
stripe-name	READ ONLY	N/A	N/A
config-dir	READ ONLY	N/A	N/A

Setting Name	Visibility	Applicability	Requirement
metadata-dir	READ ONLY	N/A	N/A
public-hostname	RUNTIME	NODE, STRIPE, CLUSTER	ALL NODES up
public-port	RUNTIME	NODE, STRIPE, CLUSTER	ALL NODES up
backup-dir	RUNTIME	NODE, STRIPE, CLUSTER	ALL NODES up
data-dirs	RUNTIME	NODE, STRIPE, CLUSTER	ALL NODES up
logger-overrides.<logger-name>	RUNTIME	NODE, STRIPE, CLUSTER	TARGETED NODES up
log-dir	RESTART	NODE, STRIPE, CLUSTER	TARGETED NODES up
client-reconnect-window	RUNTIME	CLUSTER	ACTIVES up
client-lease-duration	RUNTIME	CLUSTER	ACTIVES up
offheap-resources	RUNTIME	CLUSTER	ACTIVES up
license-file	RUNTIME	CLUSTER	ALL NODES up
cluster-name	RUNTIME	CLUSTER	ALL NODES up
failover-priority	RESTART	CLUSTER	ALL NODES up
security-dir	RESTART	CLUSTER	ALL NODES up
audit-log-dir	RESTART	CLUSTER	ALL NODES up
authc	RESTART	CLUSTER	ALL NODES up
ssl-tls	RESTART	CLUSTER	ALL NODES up
whitelist	RESTART	CLUSTER	ALL NODES up

Common options

1. -v (long option --verbose)

This option gives you a verbose output, and is useful to debug error conditions. Default: false

2. -srd (long option --security-root-directory)

This option can be used to communicate with a server which has TLS/SSL-based security configured. For more details on setting up security in a Terracotta cluster, see [“Security Core Concepts” on page 108](#).

3. `-t` (long option `--connection-timeout`)

This option lets you specify a custom timeout value for connections to be established. Default: 10s

4. `-r` (long option `--request-timeout`)

This option lets you specify a request timeout value for operations. Default: 10s

Performing cluster activation and topology changes

Activating a cluster using `activate` command

Here is the CLI to activate all the nodes of a cluster, in case the cluster is not yet CONFIGURED:

Syntax:

```
activate (-connect-to <hostname[:port]> | [-config-file <config-file>])
  [-cluster-name <cluster-name>] [-license-file <license-file>]
  [-restart-wait-time <restart-wait-time>] [-restart-delay <restart-delay>]
  [-restrict]
```

Parameters:

■ `-connect-to <hostname[:port]>`

Node to connect to

■ `-cluster-name <cluster-name>`

Cluster name

■ `-config-file`

Config file containing nodes to be activated

■ `-license-file <license-file>`

License file

■ `-restart-wait-time <restart-wait-time>`

Maximum time to wait for the nodes to restart. Default: 120s

■ `-restart-delay <restart-delay>`

Delay before the server restarts itself. Default: 2s

■ `-restrict`

Restrict the activation process to the node only

The activation process:

1. Validates the configuration consistency
2. Validates the license
3. Writes the validated cluster config inside the config directory of all nodes
4. Restarts the nodes

Once a cluster is activated, it becomes usable for Terracotta clients.

Adding nodes using the attach command

Here is the general CLI to attach/detach to update the topology:

Syntax:

```
attach (-to-cluster <hostname[:port]> -stripe <hostname[:port]> |
  -to-stripe <hostname[:port]> -node <hostname[:port]>)
  [-force] [-restart-wait-time <restart-wait-time>]
  [-restart-delay <restart-delay>]
```

Parameters:

- `-stripe <hostname[:port]>`
Stripe to add
- `-to-cluster <hostname[:port]>`
Destination cluster
- `-node <hostname[:port]>`
Node to add
- `-to-stripe <hostname[:port]>`
Destination stripe
- `-restart-delay <restart-delay>`
Delay before the server restarts itself
- `-restart-wait-time <restart-wait-time>`
Maximum time to wait for the nodes to restart
- `-force`
Force the operation

Examples:

1. Attaching 2 UNCONFIGURED nodes to form a single stripe:

```
# We start an UNCONFIGURED node
> start-tc-server.sh -name=node-1 -y availability -hostname=10.0.0.1
-config-dir=/path/to/node-1-config-dir
# we start another UNCONFIGURED node
> start-tc-server.sh -name=node-2 -y availability -hostname=10.0.0.2
-config-dir=/path/to/node-2-config-dir
# we attach node 2 in the same stripe as node 1
> config-tool.sh attach -to-stripe 10.0.0.1 -node 10.0.0.2
```

2. Attaching a node to a stripe:

```
# Continuing from the previous example, we start a third UNCONFIGURED node
> start-tc-server.sh -name=node-3 -y availability -hostname=10.0.0.3
-config-dir=/path/to/node-3-config-dir
# we attach node 3 in the same stripe as node 1
> config-tool.sh attach -to-stripe 10.0.0.1 -node 10.0.0.3
```

3. Attaching a stripe to a cluster:

```
# Continuing from the previous example, we start 3 more UNCONFIGURED nodes
> start-tc-server.sh -name=node-4 -y availability -hostname=10.0.0.4
-config-dir=/path/to/node-4-config-dir
> start-tc-server.sh -name=node-5 -y availability -hostname=10.0.0.5
-config-dir=/path/to/node-5-config-dir
> start-tc-server.sh -name=node-6 -y availability -hostname=10.0.0.6
-config-dir=/path/to/node-6-config-dir
# we attach these nodes together to form a stripe
> config-tool.sh attach -to-stripe 10.0.0.4 -node 10.0.0.5
> config-tool.sh attach -to-stripe 10.0.0.4 -node 10.0.0.6
# we attach this stripe to the cluster formed in the previous example
> config-tool.sh attach -to-cluster 10.0.0.1 -stripe 10.0.0.4
```

Removing nodes using the detach command

Syntax:

```
> config-tool.sh detach (-from-cluster <hostname[:port]>
-stripe <hostname[:port]> | -from-stripe <hostname[:port]> -node <hostname[:port]>)
[-force] [-stop-wait-time <restart-wait-time>] [-stop-delay <stop-delay>]
```

Parameters:

- `-stripe <hostname[:port]>`
Stripe to remove
- `-from-cluster <hostname[:port]>`
Destination cluster
- `-node <hostname[:port]>`
Node to remove
- `-from-stripe <hostname[:port]>`
Destination stripe

- `-stop-delay <stop-delay>`
Delay before the server stops itself. Default: 2s
- `-stop-wait-time <stop-wait-time>`
Maximum time to wait for the nodes to stop. Default: 120s
- `-force`
Force the operation

Examples

1. Detaching a node from a stripe before cluster activation:

```
> config-tool.sh detach -from-stripe 10.0.0.1 -node 10.0.0.2
```

2. Detaching a stripe from a cluster before cluster activation:

```
> config-tool.sh detach -from-cluster 10.0.0.1 -stripe 10.0.0.4
```

Performing configuration changes

Retrieving cluster configuration using `get` command

Syntax:

```
get -connect-to <hostname[:port]> [-runtime]
    -setting <[namespace:]setting>,<[namespace:]setting>...
```

Parameters:

- `-connect-to <hostname[:port]>`
The node to connect to
- `-setting <[namespace:]setting>,<[namespace:]setting>...`
List of settings to get
- `-runtime`
Get the runtime value of the setting

Note:

namespace determines the scope of the update. The format can be `stripe.<stripeId>.node.<nodeId>`: to apply a change only on a specific node, or `stripe.<stripeId>`: to apply a change only on the nodes on a stripe, or no namespace to apply a change on all nodes of the cluster.

The namespace depends on the setting: some settings can only be applied cluster-wide, some per node or per stripe.

Examples:

1. Get the value of a setting from a cluster:

```
> config-tool.sh get -connect-to 10.0.0.1 -setting offheap-resources.foo  
offheap-resources=foo:512MB
```

2. Get the values of some offheap resources on a cluster:

```
> config-tool.sh get -connect-to 10.0.0.1 -setting offheap-resources.foo -setting  
offheap-resources.bar  
offheap-resources.foo=512MB  
offheap-resources.bar=128MB
```

3. Get the values of all the offheap resources on a cluster:

```
> config-tool.sh get -connect-to 10.0.0.1 -setting offheap-resources  
offheap-resources=foo:512MB,bar:128MB,baz:64MB
```

Updating settings with set command

Syntax:

```
set -connect-to <hostname[:port]> -setting  
<[namespace:]setting=value>,<[namespace:]setting=value>...
```

Parameters:

- `-connect-to <hostname[:port]>`
The node to connect to
- `-setting <[namespace:]setting=value>,<[namespace:]setting=value>...`
List of settings with their values to be set

Examples:

1. Set some settings on a cluster:

```
> config-tool.sh set -connect-to 10.0.0.1 -setting offheap-resources.foo=512MB
```

2. Set offheap-resource bar to 512MB and foo to 1GB in the cluster:

```
> config-tool.sh set -connect-to 10.0.0.1 -setting offheap-resources.foo=1GB -setting  
offheap-resources.bar=512MB
```

Note:

The CLI supports several `-setting <setting>` parameters, to set or unset several settings at once. They together form a "change-set" which is either applied fully or not applied at all.

Removing settings with the `unset` command

Syntax:

```
unset -connect-to <hostname[:port]> -setting
<[namespace:]setting>,<[namespace:]setting>...
```

Parameters:

- `-connect-to <hostname[:port]>`
The node to connect to
- `-setting <[namespace:]setting>,<[namespace:]setting>...`
List of settings to be unset

Examples:

1. Remove logger-overrides from a node:

```
> config-tool.sh unset -connect-to 10.0.0.1 -setting stripe.1.node.1:logger-overrides
```

2. Remove backup-dir from a cluster:

```
> config-tool.sh unset -connect-to 10.0.0.1 -setting backup-dir
```

Exporting cluster configuration

Syntax:

```
export -connect-to <hostname[:port]> [-output-file <config-file>] [-include-defaults]
[-runtime]
```

Parameters:

- `-connect-to <hostname[:port]>`
The target cluster to export
- `-output-file <config-file>`
The output config file. If this option is not specified, the configuration is displayed on the console.
- `-include-defaults`
Include the settings having system default values as well. Default: false
- `-runtime`
Export the runtime config instead of the on-disk config. Default: false

Examples:

1. Export the configuration of a single stripe cluster containing two nodes, and display it on the console:

```
> config-tool.sh export -connect-to 10.0.0.1
cluster-name=my-cluster
failover-priority=consistency:3
authc=file
ssl-tls=true
stripe.1.node.1.hostname=1.company.internal
stripe.1.node.1.name=node-1
stripe.1.node.1.audit-log-dir=%H/terracotta/audit
stripe.1.node.1.security-dir=%H/terracotta/security
stripe.1.node.2.hostname=2.company.internal
stripe.1.node.2.name=node-2
stripe.1.node.2.port=9510
stripe.1.node.2.group-port=9530
stripe.1.node.2.audit-log-dir=%H/terracotta/audit
stripe.1.node.2.security-dir=%H/terracotta/security
```

2. Export the configuration of a cluster containing node with address 10.0.0.1, and save it in a file:

```
> config-tool.sh export -connect-to 10.0.0.1
-output-file=/path/to/destination/config-file
```

Importing cluster configuration

The cluster configuration in config file format can be imported on to running nodes of an unconfigured cluster to define its topology and configuration.

Syntax:

```
> import -config-file <config-file> [-node <hostname[:port]>]
```

Parameters:

- `-config-file <config-file>`

The config file to import the configuration from

- `-node <hostname[:port]>`

Node to connect to

Diagnosing and fixing problems

Viewing cluster health using "diagnostic" command

The diagnostic command displays the online, activation, health, restart, state and configuration records for each node in the cluster.

Syntax:

```
diagnostic -connect-to <hostname[:port]>
```

Parameters:

- `-connect-to <hostname[:port]>`

The node to connect to

Examples:

The following example shows the execution of the `diagnostic` command on a single-stripe cluster containing two nodes, with one node activated and online, and the other node offline:

```
Diagnostic result:
[Cluster]
- Nodes: 2 (node-1@10.0.0.1:9410, node-2@10.0.0.2:9510)
- Nodes online: 1 (node-1@10.0.0.1:9410)
- Nodes online, configured and activated: 1 (node-1@10.0.0.1:9410)
- Nodes online, configured and in repair: 0
- Nodes online, new and being configured: 0
- Nodes pending restart: 0
- Configuration state: The cluster configuration is healthy. New configuration changes
  are possible.
- Configuration checkpoint found across all online configured nodes (activated or in
  repair): YES (Version: 1, UUID: e9e0b92c-db10-4275-842e-a8bf6210e97b, At:
  2020-09-03T13:06:07.935, Details: Activating cluster: my-cluster)
[node-1@10.0.0.1:9410]
- Node state: ACTIVE
- Node online, configured and activated: YES
- Node online, configured and in repair: NO
- Node online, new and being configured: NO
- Node restart required: NO
- Node configuration change in progress: NO
- Node can accept new changes: YES
- Node current configuration version: 1
- Node highest configuration version: 1
- Node last configuration change UUID: e9e0b92c-db10-4275-842e-a8bf6210e97b
- Node last configuration state: COMMITTED
- Node last configuration created at: 2020-09-03T13:06:07.935
- Node last configuration created from: admin-machine
- Node last configuration created by: admin
- Node last configuration change details: Activating cluster: my-cluster
- Node last mutation at: 2020-09-03T13:06:08.035
- Node last mutation from: admin-machine
- Node last mutation by: admin
[node-2@10.0.0.2:9510]
- Node state: UNREACHABLE
- Node online, configured and activated: NO
- Node online, configured and in repair: NO
- Node online, new and being configured: NO
```

Fixing cluster health using "repair" command

The `repair` command can be used to fix cluster configuration inconsistency problems. The result of running this command on a healthy cluster is a no-op.

Important:

Please contact Software AG support before running this command

Syntax:

```
repair -connect-to <hostname[:port]> [-force commit|rollback|reset]
```

Parameters:

- `-connect-to <hostname[:port]>`
The node to connect to
- `-force commit|rollback|reset`
Force a commit, rollback or reset operation

13 Parameter Substitution

Parameter substitution provides a way to substitute variables with pre-defined system properties in the Terracotta Server configuration file. Thus, a significant number of fields can be intelligently populated based on machine specific properties. Parameter substitution is commonly done for hostnames, IP addresses and directory paths.

The following predefined substitutions are available for use:

Parameter	Description
%h	the fully-qualified host name of the machine
%i	the IP address of the machine corresponding to localhost
%D	the time stamp corresponding to the current date-time in yyyyMMddHHmmssSSS format
%H	the user's home directory corresponding to the user.home Java system property
%n	the username corresponding to the user.name Java system property
%o	the operating system name corresponding to the os.name Java system property
%a	the processor architecture corresponding to the os.arch Java system property
%v	the operating system version corresponding to the os.version Java system property
%t	the temporary directory corresponding to the java.io.tmpdir Java system property
%d	unique temporary directory
%(system property)	a standard or custom Java system property. If a custom Java property needs to be used, it should be first set in the JVM by setting it in JAVA_OPTS in the -Djava.custom.property=value format.

Note:

The variable `%i` is expanded into a value determined by the host's networking setup. In many cases that setup is in a `hosts` file containing mappings that may influence the value of `%i`. Test this variable in your production environment to check the value it interpolates.

14 Configuring the Terracotta Server

Overview

For your application end-points to be useful they must be able to utilize storage resources configured in your Terracotta Servers. The services offered make use of your server's underlying JVM and OS resources, including direct-memory (offheap) and disk persistence.

A Note on System Requirements and Considerations

We recommend to familiarize yourself with the following recommendations concerning the Terracotta Server:

Planning a Successful Deployment (*Installation Guide*)

System Recommendations for Hybrid Caching (*in Terracotta Server Administration Guide > Configuring the Terracotta Server*)

System Recommendations for Fast Restart (FRS) (*in Terracotta Server Administration Guide > Configuring the Terracotta Server*)

Offheap Resources

The use of JVM Direct-Memory (offheap) is a central part of the operation of a Terracotta Server. In effect, you **must** allocate and make available to your server enough offheap memory for the proper operation of your application.

In your config file you define one or more named *offheap resources* of a fixed size. These named resources are then referred to in the configuration of your application end-points to allow for their usage.

Refer to the section *Clustered Caches* in the *Ehcache API Developer Guide* for more details about the use of offheap resources.

Offheap Resource Configuration

Offheap resources can be configured with the `offheap-resources` property. By default, an offheap resource with name `main` and size `512MB` is defined on the server.

The following snippet shows the configuration of two offheap resources - `primary-server-resource` with size 384MB, and `secondary-server-resource` with size 256MB.

```
offheap-resources=primary-server-resource:384MB,secondary-server-resource:256MB
```

Data Directories

A data directory is a location on disk, identified by a name, and mapped to a disk location, where a Terracotta Server's data resides.

Data directories are commonly configured by server administrators and specified in the Terracotta Server configuration. Data directory names can be used by products that need durable storage for persistence and fast restart from crashes. For example, restartable cache managers need to be supplied with a data directory name to persist the restartable `CacheManager` specific data.

For information on restartable servers, see the section [“Server Persistence” on page 60](#) below. See also the sections *Fast Restartability* and *Creating a Restartable Cache Manager* in the *Ehcache API Developer Guide*.

Data Directories Configuration

Data directories can be configured with the `data-dirs` property. By default, a data directory with name `main` is defined which maps to disk location `%(user.home)/terracotta/user-data/main`.

The following snippet shows the configuration of two data directories - `someData` with disk location `/mnt1/data`, and `otherData` with location `%(logs.path)/data`.

```
data-dirs=someData:/mnt1/data,otherData:%(logs.path)/data
```

`%(logs.path)` gets substituted with the `logs.path` system property. Visit the section [“Parameter Substitution” on page 57](#) for the complete list of substitutable parameters.

General Notes on Configuring Data Directories

- A data directory specified on a server must be specified on all the servers in the cluster.
- Each data directory must be given a unique mount point (or disk location).
- The data directories are created if they do not exist already.
- Changing the disk location of the data directory between server restarts, without copying the data, is equivalent to erasing that data. It will cause unpredictable runtime errors that depend on the exact data lost.

Server Persistence

The Terracotta server saves its internal state on a disk which enables server restarts without loss of data.

Care must be taken to avoid losing data when restarting the stripe. Refer to the section [“Restarting a Stripe” on page 101](#) for more details. Passive restartable servers automatically back up their data at restart for safety reasons. Refer to the topic *Passive servers* in the section [“Active and Passive Servers” on page 9](#) for more details.

Server Persistence Configuration

Server persistence is configured with the `metadata-dir` property. By default, it is set to `%(user.home)/terracotta/metadata`.

The following snippet shows the configuration of `metadata-dir` to disk location `/path/to/metadata-dir`:

```
metadata-dir=/path/to/metadata-dir
```

Relation to Fast Restartability

The Ehcache *Fast Restartability* feature depends on, and makes use of, server persistence.

Refer to the section *Fast Restartability* in the *Ehcache API Developer Guide* for more information.

15 System Recommendations for Hybrid Caching

Hybrid Caching supports writing to one single mount, so all of the Hybrid capacity must be presented to the Terracotta process as one continuous region, which can be a single device or a RAID.

The mount should be used exclusively for the Terracotta server process. The software was designed for usage on local drives (SSD/Flash in particular) - SAN/NAS storage is not recommended. If you utilize SAN/NAS storage you will experience notably reduced and inconsistent performance - any support requests related to performance or stability on such deployments will require the user to reproduce the issue with local disks.

Note:

System utilization is higher when using Hybrid Caching, and it is not recommended to run multiple servers on the same machine. Doing so could result in health checkers timing out, and killing or restarting servers. Therefore, it is important to provision sufficient hardware, and it is highly recommended to deploy servers on different machines.

Hybrid Caching is described in detail in the Developer Guide.

16 System Recommendations for Fast Restart (FRS)

Fast Restart (FRS) supports writing to one single mount, which can be a single device or a RAID.

The mount should be used exclusively for the Terracotta server process. The software was designed for usage on local drives (SSD/Flash in particular) - SAN/NAS storage is not recommended. If you utilize SAN/NAS storage you will experience notably reduced and inconsistent performance - any support requests related to performance or stability on such deployments will require the user to reproduce the issue with local disks.

Fast Restartability is described in detail in the Developer Guide.

17 Failover Tuning

Overview

In a clustered environment, network, hardware, or other failures can cause an active server to get disconnected from the rest of the servers in its stripe. When your cluster needs to remain tolerant to such failures, you have a choice to make: choose either *consistency* or *availability* but not both (CAP theorem). If consistency is chosen over availability, then the cluster will halt processing client requests as consistent reads/writes can't be guaranteed when the cluster is partitioned. But when availability is chosen over consistency, the cluster will respond to client requests even when the cluster is partitioned but the response is not guaranteed to be consistent. In the absence of such failures, the cluster can provide both consistency and availability.

If the cluster is tuned to favor availability over consistency, then when such failures happen, the behavior of a stripe is that the remaining passive servers will then run an election and, if not able to find the old active server, the passive server that wins the election becomes the new active server. While this configuration ensures high availability of the data, risks of experiencing a so-called split-brain situation during such elections are increased. In the case of a TSA, split-brain would be a situation in which multiple servers in a stripe are acting as active servers. For example, if an active server gets partitioned from its peers in that stripe, the active server will remain active and the passive servers on the other side of the partition would elect a new active server as well. Any further operations performed on the data are likely to result in inconsistencies.

When tuned for consistency, a stripe would need at least a majority of servers connected with each other to elect an active server. Thus, even if the stripe gets partitioned into two sets of servers due to some network failure, the set with the majority of servers will elect an active server among them and proceed. In the absence of a majority, an active server will not be elected and hence the clients will be prevented from performing any operations, thereby preserving data consistency by sacrificing availability.

Server configuration

When configuring the stripe, the user needs to choose between availability and consistency as the failover priority of the stripe. To prevent split-brain scenarios and thereby preserve data consistency, failover priority must be set to consistency. However, if availability is preferred, `failover-priority` can be set to availability at the risk of running into split-brain scenarios. The following snippet shows how to configure a stripe for consistency:

Configured via config file:

```
failover-priority=consistency
```

Configured via command line during server startup (use `-y` or `-failover-priority` option):

```
start-tc-server.sh -y consistency
```

Similarly, the stripe can be tuned for *availability* as follows:

```
failover-priority=availability
```

Note:

`failover-priority` is a mandatory parameter that must be provided during server startup.

Choosing Consistency versus Availability

A Terracotta Server Array (TSA), being a distributed system, is subject to the constraints of the *CAP Theorem*. The CAP Theorem states that it is impossible for a distributed system to simultaneously provide guarantees for Consistency, Availability, and Partition tolerance. A TSA always seeks to be tolerant of network partitions so a choice must be made between data consistency and service availability.

For a TSA supporting only Ehcache users, choosing availability over consistency can be an effective and proper choice. If a network partition occurs leaving two servers unable to communicate with each other (to coordinate data) but each able to serve clients, choosing availability allows continuing operations using both servers running independently. Data residing on the separated servers will no longer be coordinated (consistency is abandoned) and will drift apart. Once the network partition is resolved and the servers are able to communicate with each other again, consistency rules are re-applied and one server is chosen as the holder of the *current* data - *the data on the other servers is discarded*. For a caching application, this may result in some cache misses and delays in some processing but, since a cache is not the system of record for the data being processed, no data is lost.

When using a TSA configured for availability over consistency to support TCStore users, data loss is a real possibility - if a network partition occurs and clients are actively updating TSA servers on both sides of the partition, data loss *will* occur once the network partition is resolved. If the applications using TCStore are tolerant of missing/inconsistent data - not an easy task - then configuring for availability over consistency is appropriate for the TSA. However, if a TCStore dataset is used as the system of record, this data loss or other inconsistency is generally undesirable if not catastrophic. If the TSA is used for a TCStore dataset which is either a system of record or for which loss/inconsistency is an undesirable outcome, then the TSA **must** be configured for consistency over availability.

You may use a single TSA supporting both Ehcache and TCStore but the choice of availability versus consistency is a cluster-level configuration - both Ehcache and TCStore users in a TSA are subject to that configuration. If your Ehcache users need high availability and your TCStore users need data consistency, you must use a separate TSA for each user base.

External voter for two-server stripes

Mandating a majority for active server election in certain topologies introduces additional availability issues. For example, in a two-server stripe the majority quorum is two as well. This

means that if these servers get disconnected from each other due to a network partition or because of a server failure, the surviving server would not promote itself as the active server as it requires 2 votes to win the election. But since the other voting server is not reachable, it will not be able to get that second vote and hence will not promote itself. In the absence of an active server, the stripe is not available.

Adding a third server is the best option, so that even if one fails, there is a majority (2 out of 3) surviving to elect an active. A three-server stripe can provide data redundancy and high availability at the same time even when one server fails. If adding a third server is not feasible, the alternate option is to get high availability without risking data consistency (via split-brain scenarios) using an external voter. But this configuration cannot offer data redundancy (like a three-server stripe) if a server fails.

An external voter is a client that is allowed to cast a vote in the election of a new active server, in cases where a majority of servers in a stripe are unable to reach a consensus on electing a new active server.

External voter configuration

The number of external voters needs to be described in the server configuration. It is recommended that the total number of servers and external voters be kept as an odd number.

External voters need to get registered with the servers to get added as voting members in their elections. If there are n voters configured in the server, then the first n voting clients requesting to get registered will be added as voters. Registration requests of other clients will be declined and put on hold until one of the registered voters gets de-registered.

Voters can de-register themselves from the cluster so that the voting rights can be transferred to other clients waiting to get registered, if there are any. A voting client can de-register itself by using APIs or by getting disconnected from the cluster.

When a voting client gets disconnected from the server, it will automatically get de-registered by the server. When the client reconnects, it will only get registered again as a voter if another voter has not taken its place while this client was disconnected.

Server configuration

A maximum count for the number of external voters allowed can optionally be added to the `failover-priority` configuration if the stripe is tuned for consistency, as follows:

```
failover-priority=consistency:3
```

1	Here you are restricting the total number of voting clients to three.
---	---

The failover priority setting and the specified maximum number of external voters across the stripes must be consistent and will be validated during the cluster configuration step. For more information on how to activate a cluster, see the section [“Performing cluster activation and topology changes” on page 48](#).

Client configuration

External voters can be of two variants:

1. Standalone voter
2. Clients using the voter library (client voter)

Standalone voter

An external voter can be run as a standalone process using a script provided with the kit. The voter script, which can be found under `tools/voter/bin/` under the product installation directory, takes the `<host>:<port>` of the machines as arguments. Each `-s` option argument must be a comma separated list of `<host>:<port>` combinations of servers in a single stripe. To register a multi-stripe cluster, multiple `-s` options can be provided for each stripe.

Usage:

```
start-tc-voter.(sh|bat) -s HOST:PORT[,HOST:PORT]... [-s HOST:PORT[,HOST:PORT]...]...
```

To connect the voter to a secure cluster, the path to the security root directory will also have to be provided using the `-srd` option. For more details on setting up security in a Terracotta cluster, see [“SSL / TLS Security Configuration in Terracotta” on page 107](#).

Client voter

Any TCStore or Ehcache client can act as an external voter as well by using a voter library distributed with the kit. A client can join the cluster as a voter by creating a `TCVoter` instance and registering itself with the cluster.

Note:

A cluster must be activated using the config tool before a client voter can be registered with it.

When the voter is no longer required, it can be de-registered from the cluster either by disconnecting that client, or by using the `deregister` API.

```
TCVoter voter = new EnterpriseTCVoterImpl(); // 1
voter.register("my-cluster-0" // 2
              "<host>:<port>","<host>:<port>"); // 3
...
voter.deregister("my-cluster-0") // 4
```

1	Instantiate a <code>TCVoter</code> instance
2	Register the voter with a cluster by providing a cluster name ...
3	and host port combinations of all servers in the cluster.
4	De-register from the cluster using the same cluster name that was used to register it.

To connect to a secure cluster, the voter must be instantiated using the overloaded constructor of `EnterpriseTCVoterImpl` that takes in the security root directory path.

Manual promotion with override voter

Since an external voter is just another process, there is no guarantee that it will always be up and available. Especially in the form of client voters, the moment the client leaves, the external voter

leaves too. In the rare event of a failure happening (partition splitting the active and passive servers or the active server crashing) and the external voter not being around either, none of the surviving servers will be acting as an active server. The servers will be stuck in a *suspended* state where operations from the regular clients are all stalled. A manual intervention will be required to get the cluster out of this state by fixing the cause of the partition or by restarting the crashed server. If neither is feasible, then the third option is to get a server manually promoted using an override vote from an external voter.

The voter process can be started in an *override* mode to promote a single server stuck in that intermediate state to be an active server. When the voter process is started in this special mode, it will connect to the server that you want to promote, give it an override vote and exit. The voter process can be started in *override* mode as follows:

```
start-tc-voter.(sh|bat) -o HOST:PORT
```

Running this command will forcibly promote the server at `HOST:PORT` to be an active server, if it is stuck in that intermediate state.

Note:

This override voting will work even if external voters are not configured in the server configuration.

Warning:

Be cautious not to start two different override voters on both sides of the partition separately so that both sides win and cause a split-brain.

Server startup

When the failover priority of the stripes is tuned for consistency, it has an impact on server startup as well. In a multi-server stripe, when the servers are started up fresh, a server will not get elected as an active server until it gets votes from all of its peers. This will require all the servers of that stripe to be brought up. Bringing up regular voters is not going to help as they need to communicate with all the active servers in the cluster to get registered. But if bringing up the other servers is not feasible for some reason, then an override voter can be used to forcibly promote that server.

18 Connection Leasing

Why Leasing

When a client carries out a write with IMMEDIATE or STRONG consistency, the server ensures that every client that could be caching the old value is informed, and the write will not complete until the server can ensure that clients will no longer serve a stale value.

Where network disruptions prevent the server communicating with a client in a timely manner, the server will close that client's connection to allow the write to progress.

To achieve this, each client maintains a lease on its connections to the cluster. If a client's lease expires, the server may decide to close that client's connection. A client may also close the connection if it realises that its lease has expired.

Lease Duration

When selecting the duration of lease, consider the range of possible client to server round-trip latencies over a network connection that can be considered as functional. The lease should be longer than the largest possible such latency.

On a server that is heavily loaded, there may be some additional delay in processing a client's request for a lease to be extended. Such a delay should be added into the round-trip network latency.

In addition, leases are not renewed as soon as they are issued, instead the client waits until some portion of the lease has passed before renewing. A guideline suitable for the current implementation is that leases should be approximately 50% longer to allow for this.

Setting long leases, however, has the downside that, when clients are unreachable by a server, IMMEDIATE writes could block for up to the duration of a lease.

The default duration of lease is currently 150 seconds.

Lease Configuration

To configure the lease duration, use the `client-lease-duration` property in the config file, or during server startup.

19 Cluster Tool

The cluster tool is a command-line utility that allows administrators of the Terracotta Server Array to perform a variety of cluster management tasks. For example, the cluster tool can be used to:

- Obtain the running status of servers
- Dump the state of running servers
- Take backups of running servers
- Promote a suspended server on startup or failover
- Shut down an entire cluster
- Perform a conditional partial shutdown of a cluster having one or more passive servers configured for high availability (for upgrades etc.)

The cluster tool script is located in `tools/bin` under the product installation directory as `cluster-tool.bat` for Windows platforms, and as `cluster-tool.sh` for Unix/Linux.

Cluster Tool commands

The cluster tool provides several commands. To list them and their respective options, run `cluster-tool.sh` (or `cluster-tool.bat` on Windows) without any arguments, or use the option `-h` (long option `--help`).

The following section provides a list of options common to all commands, and thus need to be specified before the command name:

Precursor options

1. `-v` (long option `--verbose`)

This option gives you a verbose output, and is useful to debug error conditions. Default: `false`.

2. `-srd` (long option `--security-root-directory`)

This option can be used to communicate with a server which has TLS/SSL-based security configured. For more details on setting up security in a Terracotta cluster, see the section [“Security Core Concepts” on page 108](#).

3. `-t` (long option `--connection-timeout`)

This option lets you specify a custom timeout value (in milliseconds) for connections to be established in cluster tool commands. Default: 10s.

4. -r (long option --request-timeout)

This option lets you specify a request timeout value for operations. Default: 10s.

The "status" Command

The status command displays the status of a cluster, or particular server(s) in the same or different clusters..

Syntax:

```
status [-n <cluster-name>] [-o json]
-s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- -n <cluster-name>

The name of the configured cluster.

- -o json

Output in JSON format. Default is tabular.

- -s <hostname[:port]>,<hostname[:port]>...

The hostname:port(s) or hostname s) (default port being 9410) of running servers, each specified using the -s option. When provided with option -n, a reachable server in the provided list will be used. Otherwise, the command will be individually executed on each server in the list.

Examples

- The example below shows the execution of a cluster-level status command.

```
./cluster-tool.sh status -n tc-cluster -s localhost
| STRIPE: 1 |
+-----+-----+-----+
| Server Name | Host:Port | Status |
+-----+-----+-----+
| server-1    | localhost:9410 | ACTIVE |
| server-2    | localhost:9510 | PASSIVE |
+-----+-----+-----+
| STRIPE: 2 |
+-----+-----+-----+
| Server Name | Host:Port | Status |
+-----+-----+-----+
| server-3    | localhost:9610 | ACTIVE |
| server-4    | localhost:9710 | PASSIVE |
+-----+-----+-----+
```

- The example below shows the execution of a server-level status command. No server is running at localhost:9510, hence the UNREACHABLE status.

```
./cluster-tool.sh status -s localhost:9410 -s localhost:9510 -s localhost:9910
```

Host-Port Additional Information	Status	Member of Cluster
localhost:9410 -	ACTIVE	tc-cluster
localhost:9510 -	PASSIVE	tc-cluster
localhost:9910 localhost:9910=Connection refused;	UNREACHABLE	-

Error (PARTIAL_FAILURE): Command completed with errors.

To learn more about server states, visit the section [“Logical Server States”](#) on page 13.

The "promote" command

The `promote` command can be used to promote a server stuck in a *suspended* state. For more information about suspended states, refer to the topics *Server startup* and *Manual promotion with override voter* in the section [“Failover Tuning”](#) on page 67.

Syntax:

```
promote -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option. The command will be individually executed on each server in the list.

Note:

There is no cluster-wide flavor for this command.

Examples

- The example below shows the execution of the `promote` command on a server stuck in suspended state at `localhost:9410`.

```
./cluster-tool.sh promote -s localhost
Following sub-operations were successful:
  localhost:9410: Server promotion successful
Command completed successfully.
```

- The example below shows the erroneous execution of a server-level `promote` command. The server running at `localhost:9510` is not in a suspended state to be promoted, hence the failure.

```
./cluster-tool.sh promote -s localhost:9510
Following sub-operations were unsuccessful:
  localhost:9510:
com.terracottatech.tools.clustertool.exceptions.ClusterToolException:
  Promote command failed as the server is not in a suspended state
Error (FAILURE): Command failed.
```

The "dump" Command

The `dump` command dumps the state of a cluster, or particular server(s) in the same or different clusters. The dump of each server can be found in its logs.

Syntax:

```
dump [-n <cluster-name>] -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-n <cluster-name>`

The name of the configured cluster.

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option. When provided with option `-n`, a reachable server in the provided list will be used. Otherwise, the command will be individually executed on each server in the list.

Examples

- The example below shows the execution of a cluster-level `dump` command.

```
./cluster-tool.sh dump -n tc-cluster -s localhost:9410
Contacting servers: [localhost:9410]
Using reachable server: localhost:9410 to carry out the operation
Following sub-operations were successful:
  localhost:9410: Dump successful
  localhost:9510: Dump successful
  localhost:9610: Dump successful
  localhost:9710: Dump successful
Command completed successfully.
```

- The example below shows the execution of a server-level `dump` command. No server is running at `localhost:9910`, hence the dump failure.

```
./cluster-tool.sh dump -s localhost:9410 -s localhost:9510 -s localhost:9910
Following sub-operations were successful:
  localhost:9410: Dump successful
  localhost:9510: Dump successful
Following sub-operations were unsuccessful:
localhost:9910:
  org.terracotta.diagnostic.client.connection.DiagnosticServiceProviderException:
  com.terracotta.connection.api.DetailedConnectionException:
  java.util.concurrent.TimeoutException: localhost:9910=Connection refused;
Error (PARTIAL_FAILURE): Command completed with errors.
```

The "ipwhitelist-reload" Command

The `ipwhitelist-reload` command reloads the IP whitelist on a cluster, or particular server(s) in the same or different clusters. See the section [“IP Whitelisting” on page 115](#) for details.

Syntax:

```
ipwhitelist-reload [-n <cluster-name>] -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-n <cluster-name>`

The name of the configured cluster.

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option. When provided with option `-n`, a reachable server in the provided list will be used. Otherwise, the command will be individually executed on each server in the list.

Examples

- The example below shows the execution of a cluster-level `ipwhitelist-reload` command.

```
./cluster-tool.sh ipwhitelist-reload -n tc-cluster -s localhost
Contacting servers: [localhost:9410]
Using reachable server: localhost:9410 to carry out the operation
Following sub-operations were successful:
  localhost:9410: IP whitelist reload successful
  localhost:9510: IP whitelist reload successful
  localhost:9610: IP whitelist reload successful
  localhost:9710: IP whitelist reload successful
Command completed successfully.
```

- The example below shows the execution of a server-level `ipwhitelist-reload` command. No server is running at `localhost:9510`, hence the IP whitelist reload failure.

```
./cluster-tool.sh ipwhitelist-reload -s localhost:9410 -s localhost:9510 -s
localhost:9910
Following sub-operations were successful:
  localhost:9410: IP whitelist reload successful
  localhost:9510: IP whitelist reload successful
Following sub-operations were unsuccessful:
  localhost:9910:
    org.terracotta.diagnostic.client.connection.DiagnosticServiceProviderException:
    com.terracotta.connection.api.DetailedConnectionException:
    java.util.concurrent.TimeoutException: localhost:9910=Connection refused;
Error (PARTIAL_FAILURE): Command completed with errors.
```

The "backup" Command

The backup command takes a backup of the running Terracotta cluster. The backup is taken on active servers only. Before taking backup of a cluster, `backup-dir` needs to be set on each server. For more details about this feature, see [“Backup, Restore and Data Migration” on page 91](#).

Syntax:

```
backup -n <cluster-name> -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-n <cluster-name>`

The name of the configured cluster.

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option. A reachable server in the provided server list will be used for connection.

Note:

There's no server-level flavor of this command, as backup works at the cluster level only.

Examples

- The example below shows the execution of a cluster-level successful backup command. Note that the server at `localhost:9610` was unreachable.

```
./cluster-tool.sh backup -n tc-cluster -s localhost:9710 -s localhost:9410
Contacting servers: [localhost:9710, localhost:9410]
Following sub-operations were unsuccessful:
  localhost:9710:
    org.terracotta.diagnostic.client.connection.DiagnosticServiceProviderException:
      com.terracotta.connection.api.DetailedConnectionException:
        java.util.concurrent.TimeoutException: localhost:9710=Connection refused;
Using reachable server: localhost:9410 to carry out the operation
PHASE 0: SETTING BACKUP NAME TO : 996e7e7a-5c67-49d0-905e-645365c5fe28
localhost:9710: TIMEOUT
localhost:9410: SUCCESS
localhost:9510: SUCCESS
localhost:9610: SUCCESS
PHASE (1/4): PREPARE_FOR_BACKUP
localhost:9710: TIMEOUT
localhost:9410: SUCCESS
localhost:9510: NOOP
localhost:9610: SUCCESS
PHASE (2/4): ENTER_ONLINE_BACKUP_MODE
localhost:9410: SUCCESS
localhost:9610: SUCCESS
PHASE (3/4): START_BACKUP
localhost:9410: SUCCESS
localhost:9610: SUCCESS
PHASE (4/4): EXIT_ONLINE_BACKUP_MODE
localhost:9410: SUCCESS
localhost:9610: SUCCESS
Command completed successfully.
```

- The example below shows the execution of a cluster-level failed backup command.

```
./cluster-tool.sh backup -n tc-cluster -s localhost:9610
Contacting servers: [localhost:9610]
Using reachable server: localhost:9610 to carry out the operation
PHASE 0: SETTING BACKUP NAME TO : 93cdb93d-ad7c-42aa-9479-6efbdd452302
localhost:9410: SUCCESS
localhost:9510: SUCCESS
localhost:9610: SUCCESS
localhost:9710: SUCCESS
PHASE (1/4): PREPARE_FOR_BACKUP
localhost:9410: SUCCESS
localhost:9510: NOOP
localhost:9610: SUCCESS
```

```
localhost:9710: NOOP
PHASE (2/4): ENTER_ONLINE_BACKUP_MODE
localhost:9410: BACKUP_FAILURE
localhost:9610: SUCCESS
PHASE (CLEANUP): ABORT_BACKUP
localhost:9410: SUCCESS
localhost:9610: SUCCESS
Error (FAILURE): Unable to complete backup.
```

The "shutdown" Command

The shutdown command shuts down a running Terracotta cluster. During the course of the shutdown process, it ensures that:

- Shutdown safety checks are performed on all the servers. Exactly what safety checks are performed will depend on the specified options and is explained in detail later in this section.
- All data is persisted to eliminate data loss.
- All passive servers are shut down first before shutting down the active servers.

The shutdown command follows a multi-phase process as follows:

1. Check with all servers whether they are OK to shut down. Whether or not a server is OK to shut down will depend on the specified shutdown options and the state of server in question.
2. If all servers agree to the shutdown request, all of them will be asked to prepare for the shutdown. Preparing for shutdown may include the following:
 - a. Persist all data.
 - b. Block new incoming requests. This ensures that the persisted data will be cluster-wide consistent after shutdown.
3. If all servers successfully prepare for the shutdown, a shutdown call will be issued to all the servers.

The first two steps above ensure an atomic shutdown to the extent possible as the system can be rolled back to its original state if there are any errors. In such cases, client-request processing will resume as usual after unblocking any blocked servers.

In the unlikely event of a failure in the third step above, the error message will clearly specify the servers that failed to shut down. In this case, use the `--force` option to forcefully terminate the remaining servers. If there is a network connectivity issue, the forceful shutdown may fail, and the remaining servers will have to be terminated using operating system commands.

Note:

The shutdown sequence also ensures that the data is stripe-wide consistent. Although, it is recommended that clients are shut down before attempting to shut down the Terracotta cluster.

Syntax:

```
shutdown [ -n <cluster-name> [-f | -i] ] -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-n <cluster-name>`

The name of the configured cluster.

- `-f | --force`

Forcefully shut down the cluster, even if the cluster is only partially reachable.

- `-i | --immediate`

Do an immediate shutdown of the cluster, even if clients are connected.

- `-s <hostname[:port]>, <hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option.

If the `-n` option is not specified, this command forcefully shuts down only the servers specified in the list. For clusters having stripes configured for high availability (with at least one passive server per stripe), it is recommended that you use the partial cluster shutdown commands explained in the section below, as they allow conditional shutdown, instead of using the `shutdown` variant without the `-n` option.

If the `-n` option is specified (i.e. a full cluster shutdown), this command shuts down the entire cluster. Servers in the provided list will be contacted for connectivity, and the command will then verify the cluster configuration with the given cluster name by obtaining the cluster configuration from the first reachable server. If all servers are reachable, this command checks if all servers in all the stripes are safe to shut down before proceeding with the command.

A cluster is considered to be safe to shut down provided the following are true:

- No critical operations such as backup and restore are going on.
- No Ehcache or TCStore clients are connected.
- All servers in all the stripes are reachable.

If either the `-f` or `-i` option is specified, this command works differently than above as follows:

- If the `-i` option is specified, this command proceeds with the shutdown even if clients are connected.
- If the `-f` option is specified, this command proceeds with the shutdown even if none of the conditions specified for safe shutdown above are met.

For all cases, the shutdown sequence is performed as follows:

1. Flush all data to persistent store for datasets or caches that have persistence configured.
2. Shut down all the passive servers, if any, in the cluster for all stripes.
3. Once the passive servers are shut down, issue a shutdown request to all the active servers in the cluster.

The above shutdown sequence is the cleanest way to shut down a cluster.

Examples

- The example below shows the execution of a cluster-level successful shutdown command.

```
./cluster-tool.sh shutdown -n tc-cluster -s localhost:9410
Contacting servers: [localhost:9410]
Using reachable server: localhost:9410 to carry out the operation
Shutting down cluster: tc-cluster
STEP (1/3): Preparing to shut down
STEP (2/3): Stopping all passive servers first
STEP (3/3): Stopping all active servers
Command completed successfully.
```

- The example below shows the execution of a cluster-level successful shutdown command that fails as one of the servers in the cluster was not reachable.

```
./cluster-tool.sh shutdown -n tc-cluster -s localhost:9410
Contacting servers: [localhost:9410]
Using reachable server: localhost:9410 to carry out the operation
Error (FAILURE): Timed out trying to reach the server
Detailed Error Status for Cluster `tc-cluster` :
  ServerError{host='localhost:9510', Error='Timed out trying to reach the server'}.
Unable to process safe shutdown request.
Command failed.
```

- The example below shows the execution of a cluster-level successful shutdown command with the force option. Note that one of the servers in the cluster was already down.

```
./cluster-tool.sh shutdown -f -n tc-cluster -s localhost:9410
Contacting servers: [localhost:9410]
Using reachable server: localhost:9410 to carry out the operation
Timed out trying to reach the server
Detailed Error Status for Cluster `tc-cluster` :
  ServerError{host='localhost:9510', Error='Timed out trying to reach the server'}.
Continuing forced shutdown.
Shutting down cluster: tc-cluster
STEP (1/3): Preparing to shut down
Timed out trying to reach the server
Detailed Error Status :
  ServerError{host='localhost:9510', Error='Timed out trying to reach the server'}.
Continuing forced shutdown.
STEP (2/3): Stopping all passive servers first
STEP (3/3): Stopping all active servers
Command completed successfully.
```

Partial Cluster Shutdown Commands

Partial cluster shutdown commands can be used to partially shut down nodes in the cluster without sacrificing the availability of the cluster. These commands can be used only on a cluster that is configured for redundancy with one or more passive servers per stripe. The purpose of these commands is to allow administrators to perform routine and planned administrative tasks, such as rolling upgrades, with high availability.

The following flavors of partial cluster shutdown commands are available:

- `shutdown-if-passive`

- `shutdown-if-active`
- `shutdown-all-passives`
- `shutdown-all-actives`

As a general rule, if these commands are successful, the specified servers will be shut down. If there are any errors due to which these commands abort, the state of the servers will be left intact.

From the table of server states described in [“The `status` Command” on page 76](#), the following are the different active states that a server may find itself in:

- `ACTIVE`
- `ACTIVE_RECONNECTING`
- `ACTIVE_SUSPENDED`

Note:

In the following sections, the term 'active servers' means servers in any of the active states mentioned above, unless explicitly stated otherwise.

Similarly, the following are the passive states for a server:

- `PASSIVE_SUSPENDED`
- `SYNCHRONIZING`
- `PASSIVE`

Note:

In the following sections, the term 'passive servers' means servers in any of the passive states mentioned above, unless explicitly stated otherwise.

The `shutdown-if-passive` Command

The `shutdown-if-passive` command shuts down the specified servers in the cluster, provided the following conditions are met:

- All the stripes in the cluster are functional and there is one healthy active server with no suspended active servers per stripe.
- All the servers specified in the list are passive servers.

Syntax:

```
shutdown-if-passive -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option.

Note:

There's no cluster-level flavor of this command.

Examples

- The example below shows the execution of a successful `shutdown-if-passive` command.

```
./cluster-tool.sh shutdown-if-passive -s localhost:9510
Contacting servers: [localhost:9510]
Stopping passive node(s): [localhost:9510] of cluster: tc-cluster
STEP (1/2): Preparing to shutdown
STEP (2/2): Stopping if Passive
Command completed successfully.
```

- The example below shows the execution of a failed `shutdown-if-passive` command, as it tried to shut down a server which is not a passive server.

```
./cluster-tool.sh shutdown-if-passive -s localhost:9410
Contacting servers: [localhost:9410]
Error (FAILURE): Unable to process the partial shutdown request.
One or more of the specified server(s) are not in passive state
or may not be in the same cluster
Discovered state of all servers are as follows:
Reachable Servers : 2
Stripe #: 1
Node: {localhost:9410} State: ACTIVE
Node: {localhost:9510} State: PASSIVE
Please check server logs for more details.
Command failed.
```

The "shutdown-if-active" Command

The `shutdown-if-active` command shuts down the specified servers in the cluster, provided the following conditions are met:

- All the servers specified in the list are active servers.
- All the stripes corresponding to the given servers have at least one server in 'PASSIVE' state.

Syntax:

```
shutdown-if-active -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option.

Note:

There's no cluster-level flavor of this command.

Examples

- The example below shows the execution of a successful `shutdown-if-active` command:

```
./cluster-tool.sh shutdown-if-active -s localhost:9410
Contacting servers: [localhost:9410]
Stopping active node(s): [localhost:9410] of cluster: tc-cluster
STEP (1/2): Preparing to shut down
STEP (2/2): Shut down if active server
Command completed successfully.
```

- The example below shows the execution of a failed `shutdown-if-active` command as the specified server was not an active server.

```
./cluster-tool.sh shutdown-if-active -s localhost:9510
Contacting servers: [localhost:9510]
Error (FAILURE): Unable to process the partial shutdown request.
One or more of the specified server(s) are not in active state
  or may not be in the same cluster.
Reachable Servers : 2
Stripe #: 1
Node : {localhost:9410} State : ACTIVE
Node : {localhost:9510} State : PASSIVE
Please check server logs for more details
Command failed.
```

The "shutdown-all-passives" Command

The `shutdown-all-passives` command shuts down all the passive servers in the specified cluster, provided the following is true:

- All the stripes in the cluster are functional and there is one active server in 'ACTIVE' state with no suspended active servers per stripe.

All passive servers in all the stripes of the cluster will be shut down when this command is run.

Syntax:

```
shutdown-all-passives -n <cluster-name> -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-n <cluster-name>`

The name of the configured cluster.

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option. These host(s) need not be passive servers.

Note:

There's no server-level flavor of this command, as it can be used only to shut down all the passive servers in the entire cluster.

The command shuts down all the passive servers in a multi-phase manner as follows:

1. Check with all servers whether it is safe to shut down as a passive server.
2. Flush any data that needs to be made persistent across all servers that are going down and block any further changes.
3. Issue a shutdown request to all passive servers if all passive servers succeed in step 2.
4. If any servers fail in step 2 or above, the shutdown request will fail and the state of the servers will remain intact.

Examples

- The example below shows the execution of a successful `shutdown-all-passives` command.

```
./cluster-tool.sh shutdown-all-passives -n tc-cluster -s localhost:9410
Contacting servers: [localhost:9410]
Using reachable server: localhost:9410 to carry out the operation
Stopping passive node(s): [localhost:9510] of cluster: tc-cluster
STEP (1/2): Preparing to shutdown
STEP (2/2): Stopping if Passive
Command completed successfully.
```

The "shutdown-all-actives" Command

The `shutdown-all-actives` command shuts down the active server of all stripes in the cluster, provided the following are true:

- There are no suspended active servers in the cluster.
- There is at least one passive server in 'PASSIVE' state in every stripe in the cluster.

The active server of all stripes of the cluster will be shut down when this command returns success. If the command reports an error, the state of the servers will be left intact.

Syntax:

```
shutdown-all-actives -n cluster-name -s <hostname[:port]>,<hostname[:port]>...
```

Parameters:

- `-n cluster-name`

The name of the configured cluster.

- `-s <hostname[:port]>,<hostname[:port]>...`

The `hostname:port(s)` or `hostname(s)` (default port being 9410) of running servers, each specified using the `-s` option. These host(s) need not be active servers.

Note:

There's no server-level flavor of this command as it can be used only to shut down all the active servers in the entire cluster.

The command shuts down all the active servers in a multi-phase manner as explained below:

1. Check with all servers whether they are safe to be shut down as active servers.
2. Flush any data that needs to be made persistent across all servers that are going down and block any further changes.
3. Issue a shutdown request to all active servers if they succeed in step 2.
4. If any servers fail in step 2 or above, the shutdown request will fail and the state of the servers will remain as before.

Examples

- The example below shows the execution of a successful `shutdown-all-actives` command. Note that the specified host was a passive server in this example. As the specified host is used only to connect to the cluster and obtain the correct state of all the servers in the cluster, the command successfully shuts down all the active servers in the cluster, leaving the passive servers intact.

```
./cluster-tool.bat shutdown-all-actives -n tc-cluster -s localhost:9510
Contacting servers: [localhost:9510]
Using reachable server: localhost:9510 to carry out the operation
Stopping active node(s): [localhost:9410] of cluster: tc-cluster
STEP (1/2): Preparing to shut down
STEP (2/2): Shut down if active server
Command completed successfully.
```

20 Licensing

This document describes the installation and update procedures for Terracotta Ehcache and Terracotta licenses.

Installing a license

A Terracotta license is installed on a Terracotta cluster using the config tool either:

- during the cluster activation process using the `activate` command, or
- any later desired time by using the `set` command

These commands ensure that:

- The license is a valid Software AG license.
- The license has not expired already.
- The Terracotta configuration files do not violate the license.

The following example activates a Terracotta cluster and installs the license file:

```
./config-tool.sh activate -l license.xml -n tc-cluster -s localhost:9410
Activating nodes: localhost:9410
License installation successful
Restarting nodes: localhost:9410
Node: localhost:9410 has restarted in state: ACTIVE
All nodes came back up
Command successful!
```

See the section [“Performing cluster activation and topology changes”](#) on page 48 for a detailed explanation of the command usage.

License expiration

License expiry checks are done every midnight (UTC time) to ensure that the license in use did not expire. Midnight here is the time at the start of the day, i.e. '00:00' hours. As an example, for a license which is valid till December 31, the midnight check on December 31 will pass, but the check on January 1 midnight will fail, and license will be deemed as expired. When a license expires, a warning message like the following will be logged every 30 minutes in the server logs:

```
ATTENTION!! LICENSE expired. Time since expiry 1 day(s)
```

The license must be renewed within 7 days of expiry. If it is not done, the cluster will be shut down with the following message in the server logs:

```
Shutting down the server as a new license is not installed within 7 days.
```

License renewal

If your license expires, a new license can be obtained by contacting Software AG support. The new license can then be installed using the config tool set command as follows:

```
./config-tool.sh set -c license-file=license.xml -s localhost:9410  
Connecting to: localhost:9410 (this can take time if some nodes are not reachable)  
License validation passed: configuration change(s) can be applied  
Applying new configuration change(s) to activated cluster: localhost:9410  
Command successful!
```

See the section [“Performing configuration changes”](#) on page 51 for a detailed explanation of the command usage.

21 Backup, Restore and Data Migration

■ Overview of Backup and Restore	92
■ Data Directory Structure	92
■ Online Backup	93
■ Offline Backup	94
■ Restore	95
■ Data Migration of Ehcache data	96
■ Technical Details	97

Overview of Backup and Restore

The Backup and Restore feature enables you as an administrator of a Terracotta cluster to take a backup of the cluster and restore it from the backed up data when required.

Terracotta supports two ways of taking a backup:

1. Online backup using the cluster-tool. This is the recommended method.
2. Manual offline backup

Restore and Ehcache data migration are manual offline processes.

Note:

Migration of TCStore data is currently not supported.

When a passive server starts and discovers it has data, the data is automatically backed up for safety reasons. However, this data is not cluster-wide consistent, and **must not** be used for restoration. Refer to the topic *Passive servers* in the section [“Active and Passive Servers”](#) on [page 9](#) for more information.

Terms

Backup and Restore : Taking a snapshot of the cluster data such that it can later be installed back on the same cluster, bringing it back to the initial state.

Data Migration : Taking a snapshot of the cluster data, but installing it on a *different* cluster, bringing it to the state of the original cluster. Data Migration is also desirable in cases when only Ehcache data is needed, and not the platform data.

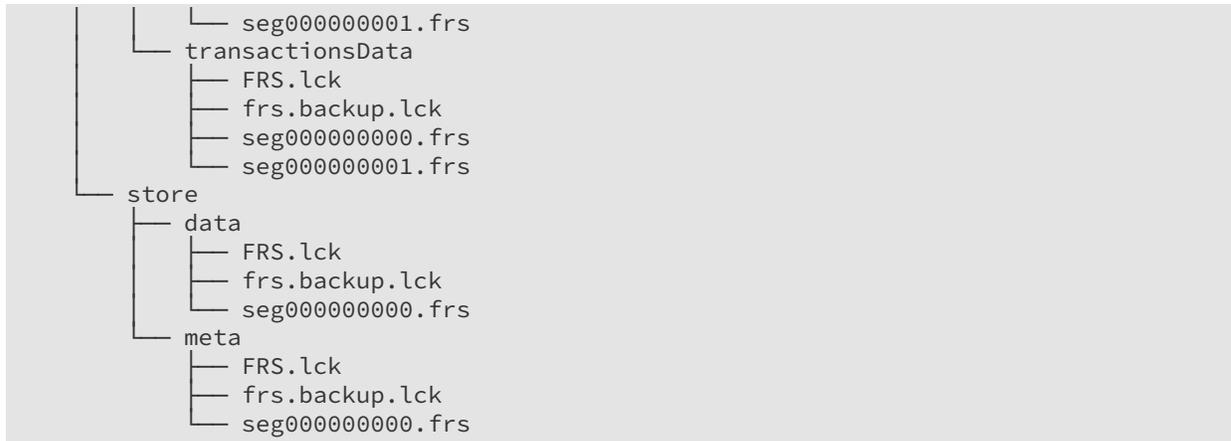
Data Directory Structure

Following is a sample data directory structure of a server containing Ehcache and TCStore data:

```

/tmp/data1/
├── server-1
│   ├── ehcache
│   │   └── frs
│   │       ├── default-frs-container
│   │       │   ├── default-cachedata
│   │       │   │   ├── FRS.lck
│   │       │   │   ├── frs.backup.lck
│   │       │   │   └── seg0000000000.frs
│   │       │   └── metadata
│   │       │       ├── FRS.lck
│   │       │       ├── frs.backup.lck
│   │       │       └── seg0000000000.frs
│   └── platform-data
│       └── entityData
│           ├── FRS.lck
│           ├── frs.backup.lck
│           └── seg0000000000.frs

```



where:

1. /tmp/data1 is the data directory path (for a given data directory) defined in the server config file
2. server-1 is the server name defined in the server config file
3. ehcache is the directory containing Ehcache data
4. platform-data is the directory containing platform specific logs
5. store is the directory containing TCStore data

Online Backup

Online backup of a Terracotta cluster is performed by the cluster-tool, and is the recommended method to take a backup. The following section describes the online backup feature and the process:

Configuring the Backup feature

To be able to take cluster backups, a backup directory must be configured. If the directory specified by the backup location path is not present, it will be created during backup. This can be done using one of the following ways:

1. Using the CLI option `-b` or `--backup-dir` during server startup
2. Using the `backup-dir` property in the config property file during server startup
3. Using the config tool `set` command to set the `backup-dir` property at runtime

Prerequisites

Before proceeding with the online backup, ensure that:

1. At least one server in each stripe is up and running.
2. The servers have read and write permissions to the backup location.
3. Backup location has enough space available to store the backup data.

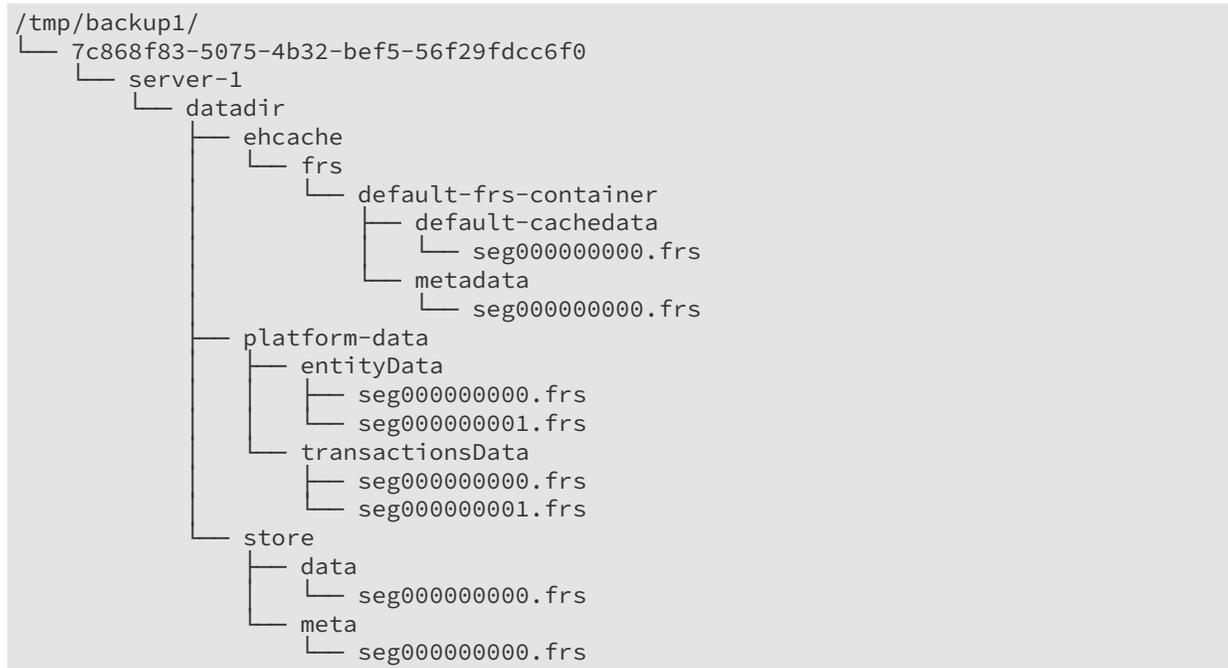
4. cluster-tool has fast connectivity to all the servers and the cluster is not heavily loaded with application requests.

Taking an online Backup

A backup is taken using the cluster-tool. Visit [“Cluster Tool” on page 75](#) for details on the backup command. If the backup fails for some reason, you can check the server logs for failure messages. Additionally, running the backup command using the `-v` (verbose) option might help.

Backup directory structure

The following diagram shows an example of the directory structure that results from a backup:



where:

1. `/tmp/backup1/` is the backup location defined in the config file
2. `7c868f83-5075-4b32-bef5-56f29fdcc6f0` is an ID created by the backup command to uniquely identify a backup instance
3. `server-1` is the server name
4. `datadir` is the data directory name (for a given data directory) defined in the server config file

Offline Backup

In the rare scenario when an online backup cannot be taken, an offline backup can be taken. The process is described as follows:

Taking an offline Backup

Follow the steps in the specified order to back up cluster data:

1. Shut down the cluster, while taking a note of the current active servers.
2. Copy the contents of the required data directories of **all** the servers which were actives prior to the shutdown to a desired location.
3. Name the directories in the manner described in the “[Backup directory structure](#)” on page 94 section above. Although this step is optional, it helps identify different instances of backup, and keeps the restore steps consistent for online and offline backup procedures.
4. Save the config files as well. These files will be used to start the stripes after a restore is performed.

Restore

The restore operation is a manual operation. During the Restore operation, you use standard operating system mechanisms to copy the complete structure (directories, subdirectories and files) of the backup into the original location. Some small structural and/or naming changes are required in the restored directories after the copy, as described in the sections below.

Note:

Restoring cache data will bring back cache entries which might have become stale by the time a restore is finished.

Performing a Restore

Before you start the Restore operation, ensure that all activity has stopped on the cluster and that the cluster is not running.

If you compare the structure of the backup under `/tmp/backup1` with the original structure under `/tmp/data1` (see both structural diagrams above), you will see some differences. You will also see that this is a single stripe cluster. Therefore, when you copy the `/tmp/backup1/<backup-name>` directory structure back to `/tmp/data1`, you need to make the following changes:

1. First choose a server as the active server for your stripe.
2. Note down the name of that server.
3. Create an empty directory for each path specified by the data directory. This will be the target directory for your restored data. Repeat this step for every data directory path specified in your config file.
4. Create a sub-directory with the name of the server under the data directories created above. For example, if the name attribute `isserver-2` for the chosen active server for this stripe and the location specified for the data directory `datadir` is `/tmp/data1`, your target directory should look like `/tmp/data1/server-2`.

5. From the backup, copy the contents of `<server-name>/<data-directory>` to this newly created directory. For example, in the example given above, copy from `/tmp/backup1/<backup-name>/server-1/data` to `/tmp/data1/server-2`
6. Start the server with the newly created data directory with the config file which was backed up from the original cluster.
7. You can now bring up the passive servers in the stripe. Please note that you don't need to copy the backup data to the passive servers as they will automatically receive the data when they synchronize with the active server. It is advisable to remove any old data on the passive servers before you bring up the passive servers.
8. Repeat the above steps for other stripes in the cluster.

Data Migration of Ehcache data

Note:

As noted above, data migration is currently not available for TCStore data.

Data migration can be performed to move Ehcache data to a new cluster without moving the platform data. Please note that only restartable caches contained in a restartable cache manager can be recovered. Since the data migration works at the data directory level, all the data of all restartable cache managers that use the same data directory will be recovered together.

How to perform an offline data migration

Follow the steps in the specified order to perform a migration of cluster data:

1. Shut down the source cluster and copy the contents of all ehcache directories from all required data directories of **all** active servers in the cluster. You can skip copying data directories containing restartable cache managers that you do not wish to migrate.
2. Start the target cluster (you can just start the active servers at this time) with the same number of stripes as the source cluster. Create the desired cache manager configuration using a client. The cluster URI (including the cluster tier manager name for the cache manager) can be different in the new cluster. If the name part of the URI is different, specify the old name as the restart identifier when using the cache manager configuration API, so that the system can map the data corresponding to a given cache manager correctly. If there are more than one cache managers under the same data directory, use the configuration API to create all the cache managers in the target cluster.

For related information, see the section *Fast Restartability* of the *Ehcache API Developer Guide*.

3. Shut down the target cluster and copy the data to the matching data directories. The data directory paths can be different on the target cluster, but must have sufficient space to contain the data being copied over.
4. Once the data is available in all the stripes, you can start the target cluster. It now loads all the cache data that was moved from the source cluster.

Technical Details

Causal and Sequential Consistency across stripes

Since TCStore and Ehcache support only causal consistency (per key) and sequential consistency (across keys for a single thread of execution), the backup image across the cluster (be it single stripe or multi-stripe) must be consistent cluster wide for the point-in-time when the backup was taken.

For instance, suppose a single thread of execution from a single client synchronously made changes to keys A, then B, then C and then D in that order. Now if the backup was captured when the client had made changes to C, intuitively the backup **MUST** have all the previous changes made to A and B, regardless of the stripe where those mutations occurred. Thus on a restore of this point-in-time backup, if the restored data has C, then it **MUST** contain the changes made to A and B. Of course, it is to be expected that such a restoration may have permanently lost D, due to the point-in-time nature of restoring from backups.

As another example, say a system had long keys from 1 to 1000 and mutated them one by one exactly in that order. If the backup had 888 as the largest key, then all keys from 1 to 887 **MUST** also exist in the backup.

Causal consistency (per key) is always implied, as a key is always within a stripe. The backup taken must be consistent for a point in time snapshot, which implies that when a snapshot is taken, all mutations/changes that happen in the system **AFTER** the snapshot is taken **MUST** not reflect in the backup.

Consistency of multiple FRS logs within a stripe

Since platform data is also backed up, there are at least two FRS logs that needs to be backed up in a consistent fashion even within a single stripe.

22 Migrating from older Terracotta versions to 10.7

A major feature of the 10.7 release is dynamic configuration of the Terracotta cluster, which changes cluster startup and configuration mechanisms. This document briefly lists the new concepts, tooling information, and the steps to be followed to migrate from an earlier 10.x to a 10.7 cluster.

Before proceeding with this document, please make sure that you have familiarized yourself with the material presented in [“Configuration Terms and Concepts”](#) on page 29.

What has changed

1. Node startup mechanism

`start-tc-server.(bat|sh)` does not support a `tc-config` XML file anymore. Several new options have been added to this script. More information can be found in the [startup script](#) section.

2. Cluster configuration format

Cluster configuration can no more be specified using `tc-config` XML files. The old `tc-configs` can be migrated to the new configuration format using [config converter tool](#).

3. Cluster configuration mechanism

The cluster tool `configure` command has been removed in favor of the [config tool](#) `activate` command. The cluster tool `reconfigure` command has been removed in favor of the `config tool set` command.

Converting old configuration files

The `config converter` tool can be used to convert from `tc-config` XML files used with Terracotta versions up to 10.5 to the new configuration format. The tool can be located under `tools/upgrade/bin` in the product installation directory, and has the following usage:

```
> config-converter.(bat|sh) convert -config <tc-config>,<tc-config>...
  ( -type directory [-license-file <license-file>] -cluster-name <new-cluster-name>
  |
  -type properties [-cluster-name <new-cluster-name>])
  [-destination <destination-dir>] [-force]
```

When the output format is `config directory` (i.e. `-type directory` or no `-type` specification), a license file can be supplied.

Note:

In a fresh install of Terracotta using the Software AG installer, the license file can be found under the Terracotta installation root.

This way, the generated config directory will be ready-for-use by Terracotta clients. This is the recommended output format. Note that there is no change in the license file, or the licensing policies. Thus, existing Terracotta licenses will continue to work with 10.7 until their expiration.

When the output format is config properties (i.e. `-type properties`), the generated config file contains information from all the `tc-config XML` files, and can also be used to start up the nodes. However, an additional cluster activation step will be required after the migrated cluster is started using this config file.

Migrating old data

If you want to use old data with 10.7 server, you need to run the create permanent entities tool. One host-port per stripe needs to be specified so that the tool can connect to all the stripes.

```
> create-permanent-entities.(bat|sh) -connect-to <hostname[:port]>,...
```

Updates to the server startup script

The `start-tc-server` script does not support `tc-config XML` files anymore. Instead, it supports several options which let you specify a node's configuration via the command itself.

See the section [“Starting and Stopping the Terracotta Server” on page 21](#) for details.

Migration steps

1. Shut down all Terracotta clients and ensure no critical operations (like backup) are running on the cluster. Note down the hosts the nodes are running on.
2. Use the cluster tool `shutdown` command to shut down the Terracotta cluster.
3. Use the [config converter tool](#) to convert `tc-config.xml` files to config directory format.
4. Copy the config directories generated from the step above to the hosts from the first step.
5. Start the nodes using the [startup script](#) with option `-r`, supplying the config directory path.
6. Use the [create-permanent-entity-tool](#) if you want to use the old data with the 10.7 server. Ensure that you run this script before connecting any clients.
7. Replace the old client jars with 10.7 jars in the client classpath.
8. Connect the clients back with the cluster.

23 Restarting a Stripe

Restart behavior is closely related to *failover*, but the difference is that the interruption period is typically much longer. On a restart, every server waits for the last active server to return instead of electing a different server as active. If a server other than the last active is elected as the active on a restart, it could cause data loss. To avoid such a data loss, every restarted server will wait in a *suspended* state until all of its peers are also started up so that the last active can be found and elected as the leader.

Unless a timeout is set, the time the clients will wait for the server to return is indefinite.

Note that a stripe can be both restartable and possess high-availability, if it is configured for restart support but also contains multiple servers. In this case, failover will progress as normal unless the entire stripe is taken offline.

Comparison with failover

The process of a *client* reconnecting to a restarted server is very similar to a newly-promoted *active* server after a fail-over. Both scenarios involve the clients reconnecting to re-send their in-flight transactions. Also, both will progress as normal once all clients have reconnected or the reconnect window closes.

The primary difference is that restart only requires one server, whereas high-availability requires at least two.

24 IPv6 support in Terracotta

This document describes the changes needed in config files, tooling, and client side APIs toward enabling IPv6 connections between clients and Terracotta cluster.

Terracotta supports IPv6 addresses as defined in the RFC 5952. Thus, all of following are acceptable:

- Full IPv6 address, e.g. `2001:db8:a0b:12f0:0:0:0:1`.
- Full IPv6 address enclosed in square brackets, e.g. `[2001:db8:a0b:12f0:0:0:0:1]`.
- Full IPv6 address enclosed in square brackets along with port, e.g. `[2001:db8:a0b:12f0:0:0:0:1]:9410`.
- Shortened IPv6 address, e.g. `2001:db8:a0b:12f0::1`.
- Shortened IPv6 address enclosed in square brackets, e.g. `[2001:db8:a0b:12f0::1]`.
- Shortened IPv6 address enclosed in square brackets along with port, e.g. `[2001:db8:a0b:12f0::1]:9410`.

Note that enclosing an IPv6 address in square brackets is mandatory only if a port is to be specified along with it.

Terracotta Server

Consider the following example:

```
stripe.1.node.1.hostname=2001:db8:a0b:12f0::1 (1)
stripe.1.node.2.hostname=[2001:db8:a0b:12f0:0:0:0:2] (2)
stripe.1.node.3.hostname=terracotta-host (3)
```

(1) Specifies a server host IP `2001:db8:a0b:12f0::1` as hostname for node 1.

(2) Specifies a server host IP `2001:db8:a0b:12f0:0:0:0:2` enclosed in square brackets as hostname for node 2.

(3) Specifies a server DNS host name `terracotta-host` which resolves to an IPv6 address as hostname for node 3.

IPv6 server sockets bind to `:::0` by default, which can be overridden using the `bind` attribute in the `<server>` element.

Command-line tools

Command-line tools which accept server addresses need to provide the IPv6 addresses of Terracotta servers as shown in the examples below:

- This example below shows the execution of the cluster tool status command for two IPv6 addresses - one enclosed in square brackets, and the other not enclosed in square brackets (and using the default port).

```
./cluster-tool.sh status -s [2001:db8:a0b:12f0::2]:9510 -s 2001:db8:a0b:12f0::1
+-----+-----+-----+
|          Host:Port          |      Status      | Member of Cluster |
+-----+-----+-----+
| [2001:db8:a0b:12f0::2]:9510 |      ACTIVE      |      tc-cluster   |
|          2001:db8:a0b:12f0::1 |      PASSIVE     |      tc-cluster   |
+-----+-----+-----+
```

Command line tools that don't accept server addresses directly do not need any change to work with IPv6.

Ehcache Client

An Ehcache client can specify the IPv6 addresses of the servers it wants to connect to either through Java APIs or XML configuration, as shown below:

API Example

```
InetSocketAddress firstServer =
    InetSocketAddress.createUnresolved("2001:db8:a0b:12f0::1", 0);
InetSocketAddress secondServer =
    InetSocketAddress.createUnresolved("2001:db8:a0b:12f0:0:0:0:2", 9510);
List<InetSocketAddress> servers = Arrays.asList(firstServer, secondServer);
String cacheManagerIdentifier = "cacheManager-1";
PersistentCacheManager cacheManager = CacheManagerBuilder
    .newCacheManagerBuilder()
    .with(EnterpriseClusteringServiceConfigurationBuilder.enterpriseCluster(servers,
        cacheManagerIdentifier) //(1)
    .autoCreate())
    .build(true);
```

1	EnterpriseClusteringServiceConfigurationBuilder.enterpriseCluster(Iterable, String) lets you create a CacheManager by specifying IPv6 addresses of the servers. The first argument is the Iterable<InetSocketAddress> of the servers in the cluster, while the second argument is the cache manager identifier.
---	---

Like other Ehcache APIs, the above API has a secure variant as well with the signature EnterpriseClusteringServiceConfigurationBuilder.enterpriseSecureCluster(Iterable, String, Path), where the last argument is the path to the client's security root directory.

XML Example

```
<ehcache:config
  xmlns:ehcache="http://www.ehcache.org/v3"
  xmlns:tc="http://www.terracottatech.com/v3/terracotta/ehcache">
  <ehcache:service>
    <tc:cluster>
      <tc:cluster-connection cluster-tier-manager="cacheManager-1"> <!-- 1 -->
        <tc:server host="[2001:db8:a0b:12f0::1]"/> <!-- 2 -->
        <tc:server host="2001:db8:a0b:12f0:0:0:0:2" port="9510" /> <!-- 3 -->
      </tc:cluster-connection>
    </tc:cluster>
  </ehcache:service>
</ehcache:config>
```

1	Cache manager identifier.
2	Terracotta server IP [2001:db8:a0b:12f0::1]. Since the port is not specified, it will default to 9410.
3	Terracotta server IP 2001:db8:a0b:12f0:0:0:0:2 and port 9510.

TCStore Client

```
InetSocketAddress firstServer =
  InetSocketAddress.createUnresolved("2001:db8:a0b:12f0::1", 0);
InetSocketAddress secondServer =
  InetSocketAddress.createUnresolved("2001:db8:a0b:12f0:0:0:0:2", 9510);
List<InetSocketAddress> servers = Arrays.asList(firstServer, secondServer);
DatasetManager datasetManager = DatasetManager.clustered(servers) // 1
  .build();
```

1	<code>DatasetManager.clustered(Iterable)</code> lets you create a <code>DatasetManager</code> by specifying IPv6 addresses of servers.
---	--

Like other TCStore APIs, the above API has a secure variant as well with the signature `DatasetManager.clustered(Iterable, Path)`, where the last argument is the path to the client's security root directory.

XML Example

```
<clustered xmlns="http://www.terracottatech.com/v3/store/clustered"> <!-- 1 -->
  <cluster-connection>
    <server host="[2001:db8:a0b:12f0::1]"/> <!-- 2 -->
    <server host="2001:db8:a0b:12f0:0:0:0:2" port="9510" /> <!-- 3 -->
  </cluster-connection>
</clustered>
```

1	Declares a clustered <code>DatasetManager</code> configuration.
2	Terracotta server IP [2001:db8:a0b:12f0::1]. Since the port is not specified, it will default to 9410.
3	Terracotta server IP 2001:db8:a0b:12f0:0:0:0:2 and port 9510.

Terracotta Management Console

To connect to a Terracotta cluster running on IPv6, you can specify the server IPv6 addresses in the **Connection URL** as shown in the snapshot:

Screenshot: Terracotta Management Console, field for Connection URL



The screenshot shows the Terracotta Management Console interface. At the top, there is a blue header bar with the text "TERRACOTTA Management Console" on the left and a hamburger menu icon on the right. Below the header, the "Connection URL:" label is followed by a text input field containing the IPv6 address "terracotta://2001:db8:a0b:12f0::1,[2001:db8:a0b:12f0:0:0:2]:9510". To the right of the input field are two buttons: "Next" and "Cancel".

25 SSL / TLS Security Configuration in Terracotta

■ Security Core Concepts	108
■ Cluster Security	113
■ TMS Security	119
■ LDAP Properties	123
■ SSL / TLS Troubleshooting guide	126

Security Core Concepts

Security features in Terracotta

Terracotta provides the following security features:

1. **Connection encryption** - encrypts client-server and server-server connections using the SSL/TLS protocol.
2. **Cluster authentication** - validates the identity of processes initiating connections to a server in a Terracotta cluster.
3. **IP whitelisting** - restricts access to only allow clients from known IP addresses.
4. **TMS authentication** - validates the identity of users attempting to use the TMS.
5. **TMS authorization** - determines if an authenticated user has access to perform a given operation in TMS.
6. **Auditing** - writes security-relevant events to audit logs.

Security Root Directory

To configure security features in Terracotta, each server and, in most cases, each client, must have a security root directory. The security root directory is a filesystem location for certificates, passwords and other security-related files.

The security root directory, and the files and subdirectories contained in it should be readable by the respective client or server process.

Important:

Ensure that the security root directory is accessible *only* by users who are permitted to run the respective Terracotta client or server.

Structure of the Security Root Directory

The directory structure below lists all possible subdirectories and files that can be present in a security root directory. Note that any one security root directory will have only a subset of these.

```
<security-root-directory>
├── access-control
│   ├── ldap.properties
│   └── users.xml
├── identity
│   ├── credentials.properties
│   └── <common-name>-<timestamp>.jks
├── trusted-authority
│   └── trusted-authority-<timestamp>.jks
└── whitelist.txt
```

Note:

It is recommended to keep only the required files and directories under the security root directory to prevent configuration errors. If any unidentifiable files or directories are found, an appropriate warning message will be logged.

Access control subdirectory

The `access-control` subdirectory should only be specified on the server-side. It contains files to configure authentication and authorization:

`ldap.properties` - configuration properties for using an LDAP server for authentication or authorization

`users.xml` - used for file-based authentication or authorization

Identity subdirectory

The `identity` subdirectory can be specified on both the client-side and the server-side. It contains certificates or password-based credentials to prove the identity of the process.

`credentials.properties` - a standard Java properties file containing two properties: `username` and `password`. These credentials are sent to the server when the connection is established, in order to authenticate. This file should be present when file-based authentication or LDAP-based authentication is configured.

`<common-name>-<timestamp>.jks` - a keystore containing a certificate for proving identity. Note that the `<common-name>` part of the filename must match the Common Name specified in the certificate and the `<timestamp>` represents the time that the certificate was created.

Trusted authority subdirectory

The `trusted-authority` subdirectory can be specified both client-side and server-side. It contains trusted root certificates.

`trusted-authority-<timestamp>.jks` - a truststore containing a trusted root certificate for validating identity certificates. Note that the `<timestamp>` represents the time that the certificate was created.

The `whitelist.txt` file

The `whitelist.txt` file should only be specified on the server-side. It contains details of client IP addresses permitted to establish connections with the cluster.

Certificates

This section assumes a good understanding of SSL/TLS fundamentals.

Certificate creation

Note:

Certificates must be created using the RSA algorithm, preferably with a key size of 4096.

The keystores and truststores, that are deployed to the `identity` and `trusted-authority` directories respectively, can be created in any way desirable as long as the following rules are followed:

Keystore creation rules

- Keystore must be of type `jks`.
- Keystore filename must be in `${common name}-${yyyyMMddThhmmss}.jks` format (e.g. `com.organization.host-20180223T102319.jks`). `yyyyMMddThhmmss` should represent the time of creation (timestamp) of the file. When multiple keystores are present, the keystore with the latest timestamp is used.
- Keystore must have only one `terracotta_security_alias` entry, and it should contain the identity certificate and the corresponding private key.
- Common Name field in the Distinguished Name in the certificate must match the common name fragment in the keystore filename. For a server the common name must match the host name.
- The identity certificate must be within its period of validity.
- The password for the keystore and the `terracotta_security_alias` store entry must be `terracotta_security_password`.
- The certificate must be created using the RSA algorithm, preferably with a key size of 4096.

Truststore creation rules

- Truststore must be of type `"jks"`.
- Truststore filename must be in `trusted-authority-${yyyyMMddThhmmss}.jks` format (e.g. `trusted-authority-20180223T102319.jks`). When multiple truststores are present, all of them are used irrespective of their timestamps.
- Truststore must have only one `terracotta_security_alias` entry, and it should contain a trusted certificate.
- The trusted certificate must be within its period of validity.
- The password for the truststore must be `terracotta_security_password`.
- The certificate must be created using the RSA algorithm, preferably with a key size of 4096.

Certificate rotation

If the certificates expire or get compromised, they must be rotated. Following are the different ways of rotating them:

Certificate rotation with cluster shutdown

The followed steps need to be performed *in order*:

1. Generate new keystore and truststore files following the rules mentioned in the *Keystore creation* and *Truststore creation* sections above.
2. Shut down all the servers and clients.
3. Replace the old keystores and truststores with the new keystores and truststores in the corresponding `identity` and `trusted-authority` directories of each client and server.
4. Start all the servers and the clients for the new certificates to take effect.

The above sequence has the advantage that it is simple to perform, with the drawback of requiring the entire cluster and the clients to be restarted.

Certificate rotation with rolling restarts

The followed steps need to be performed *in order*:

1. Generate new keystore and truststore files following the rules mentioned in the *Keystore creation* and *Truststore creation* sections above.
2. Deploy new truststores in corresponding `trusted-authority` directories of each client and server.
3. Restart all the passive servers and clients. Once the passive servers reach `PASSIVE-STANDBY` status, restart all the active servers.
4. Replace old keystores with the new keystores in corresponding `identity` directories of each client and server.
5. Restart all the passive servers and clients. Once the passive servers reach `PASSIVE-STANDBY` status, restart all the active servers.
6. Delete old truststores from corresponding `trusted-authority` directories of each client and server.
7. Restart all the passive servers and clients. Once the passive servers reach `PASSIVE-STANDBY` status, restart all the active servers.

The above sequence has the advantage that it does not require cluster downtime, with the drawback of having a significant number of steps.

Auditing

The Terracotta server can audit when security-relevant events occur. You can configure this by specifying an audit directory.

Important:

Ensure that the audit directory is accessible only by users who are permitted to run the respective Terracotta server or by users who are allowed to read the audit log.

Audit directory structure

The auditing process creates a hierarchical filesystem structure under the audit directory:

```
<audit-directory>
├── cluster-audit-logs
│   ├── <server-name>
│   │   └── <yyyy-MM-dd>
│   │       ├── terracotta-server-audit-<yyyy-MM-dd>.0.log
│   │       └── terracotta-server-audit-<yyyy-MM-dd>.1.log
├── tmc-audit-logs
│   ├── <TMC spring app name>
│   │   └── <yyyy-MM-dd>
│   │       ├── tmc-audit-<yyyy-MM-dd>.0.log
│   │       └── tmc-audit-<yyyy-MM-dd>.1.log
```

When an auditable event occurs, an appropriate message is written to the audit log file. Each line in the audit log corresponds to a single event.

Audit log files generated by a server in a Terracotta cluster are stored under the `cluster-audit-logs` directory. Audit log files generated by the TMS are stored under the `tmc-audit-logs` directory.

Each server has a separate directory named after the server name. This allows multiple servers to use the same base audit directory.

Under the server-specific directory, logs are organised by date. If any audit log file exceeds 10 MB, a new audit log file is created.

Audit entry sanitization

Audit entries are sanitized to prevent attackers rendering the audit logs illegible. The sanitization process follows the rules :

- All ASCII control characters, and extended ASCII characters are sanitized.
- All non-ASCII characters will be sanitized.
- Tilde (~), asterisk (*), pipe (|), plus (+) and backtick (`) characters are sanitized.
- Any number of occurrences of the disallowed characters, and more than one consecutive occurrence of space (` `) is sanitized.
- Input larger than 2000 characters is truncated, and a '*' is added at the end to indicate the truncation.

If the sanitization process finds no unacceptable characters or character sequences, then the information is logged surrounded with backticks. For example, if the user 'alex' successfully authenticated, that is audited as:

```
LDAP authentication success:- IP: `203.0.113.1`, User: `alex`
```

However, if the sanitization process finds unacceptable characters or character sequences, the information is logged surrounded with tildes and prefixed with `sanitized`. Characters replaced in the sanitization process appear in the format `|U+xxxx|` where the `xxxx` is the unicode codepoint

of the sanitized character in hexadecimal. For example, if the user jürgen successfully authenticated, that is audited as:

```
LDAP authentication success:- IP: `203.0.113.1`, User: ~sanitized j|U+00fc|rgen~
```

Note that there may be more than four digits in the hex codepoint due to supplementary unicode characters.

Cluster Security

Connection encryption

You can enable connection encryption using the `ssl-tls` property. When `<ssl-tls>` is specified, you must also supply:

- `security-dir` with the path to the security root directory.
- `authc` with the authentication scheme to use. See “[Authentication](#)” on page 113 for more information.

With SSL/TLS configured, the specified security root directory must contain a valid truststore in the `trusted-authority` subdirectory and a valid keystore in the `identity` subdirectory. The identity certificate is required even if file-based or LDAP-based authentication is chosen. The choice of authentication scheme may require presence of additional files in the security root directory.

On the client, the security root directory must contain a trusted authority certificate in the `trusted-authority` subdirectory. Again, the choice of authentication scheme may mean other files, including an identity certificate, are also required in the security root directory.

Authentication

To configure authentication in the cluster, use `authc` with the appropriate authentication scheme. Supported authentication schemes are `file`, `ldap`, and `certificate`. When authentication is configured, a `security-dir`, with the path to the security root directory, must also be specified.

Important:

It is highly recommended that if you configure an authentication scheme, you also configure encrypted connections using the `ssl-tls` property, otherwise an attacker could acquire credentials by eavesdropping on the unencrypted connection.

Certificate-based authentication

To configure certificate-based authentication, use the `certificate` authentication scheme in the `authc` property, and set `ssl-tls` to `true`. Also, clients must have an appropriate identity certificate keystore in their security root directory.

File-based authentication

To configure file-based authentication, use the `file` authentication scheme in the `authc` property.

The server's security root directory must contain a `users.xml` file, which is a list of all valid users with a password hash for each user. The only password hashing algorithm currently supported is `bcrypt`.

Example of a `users.xml` file for authentication:

```
<users>
  <user>
    <username>alex</username>
    <password>
      <algorithm>bcrypt</algorithm>
      <hash>$2a$10$UoM85/5I4SnIbr0QuFZ43ekffuQKSxZmL93bR9VMcdr2URmPyjyX2</hash>
    </password>
  </user>
  <user>
    <username>beth</username>
    <password>
      <algorithm>bcrypt</algorithm>
      <hash>$2a$10$6D6c79lE0k/0SxrEtnfhGe2Yr.ygG0rFP1QzeyD9qshIMRrpUMOAS</hash>
    </password>
  </user>
</users>
```

Note:

Hashes should start with `bcrypt` version `$2a$`.

Generating `bcrypt` password hashes

The command-line utility for generating `bcrypt` password hashes is located in `tools/security/bin` under the product installation directory as `bcrypt.bat` for Windows platforms, and as `bcrypt.sh` for Unix/Linux.

When running the `bcrypt` script, you must specify the number of rounds. The number depends on the performance of the server hardware and how you decide to trade off security with speed. A higher number creates hashes that are harder for attackers to crack, but are also harder for your servers to verify. Increasing the number of rounds by 1 doubles the difficulty. You may find that a value between 10 and 13 is suitable.

The `bcrypt` script can read the input from the console:

```
bcrypt.sh -n 10
```

or directly from the command line:

```
bcrypt.sh -n 10 pa$$w0rd
```

The console flavor of the command should be preferred to prevent the shell from saving the input password in its history.

LDAP-based authentication

To configure LDAP-based authentication, use the `ldap` authentication scheme in the `authc` property.

The server's security root directory must contain an `ldap.properties` file. The properties in the `ldap.properties` file are the same properties that are used to configure LDAP integration in other Software AG products. See “[LDAP Properties](#)” on page 123 for a full list of supported properties.

Example of an `ldap.properties` file for authentication:

```
url=ldap://ldapserver.example.com:389
userrootdn=ou=People,dc=example,dc=com
uidprop=uid
personobjclass=person
```

Auditing

To configure security event auditing, use the `audit-log-dir` property along with `security-dir` and at least one form of security (i.e. one of `ssl-tls`, `authc`, or `whitelist`). The directory specified in `audit-log-dir` must exist already, and must be appropriately access controlled to prevent illegitimate access to audit logs.

IP Whitelisting

Introduction

The IP whitelisting feature enables you as the cluster administrator to ensure that only clients from known IP addresses can access the TSA. You can use this feature to prevent malicious clients from establishing connections to the TSA.

Note:

It should be understood that usage of this feature in itself does not provide a strong level of security for the TSA. The ideal way to enforce connection restrictions based on IP addresses would be to use host-level firewalls.

The whitelist file

A whitelist file is a plain-text file containing a list of IPs. Only clients configured with these IPs are allowed to access the TSA. The server IPs specified in the config file, and the localhost IPs of the server are always whitelisted. The whitelist file must be named `whitelist.txt` and placed in the security root directory.

Note:

An empty whitelist file has the semantics of blacklisting all IPs, except the ones fetched from the config file, and those corresponding to `localhost`.

The whitelist file follows these parsing rules:

1. The entries can be IP addresses, or CIDR notations (to represent IP ranges). Any entry that is not a valid IP address or a valid CIDR is ignored.
2. Each line in the file can contain either a single IP address, or a comma-separated list of IP addresses.
3. Lines beginning with `#` are considered as comments, and are ignored during parsing.
4. Blank lines are ignored.

The following is an example of a valid whitelist file:

```
# Caching clients
192.168.5.28, 192.168.5.29, 192.168.5.30
10.60.98.0/28
# Other clients
192.168.10.0/24
```

Usage

To configure IP whitelisting, use `whitelist` along with the `security-dir` property.

If the `whitelist.txt` file is not found in the security root directory, or there is an error reading the file, the server startup will fail with an appropriate error message.

If hostnames are used in the config file, the server attempts to resolve these hostnames to IPs. If the resolution fails, the server startup fails with an appropriate error message. Note that hostname resolution is done for the config file only, and any hostnames present inside the `whitelist.txt` file are ignored.

A multi-stripe cluster should be started with the same `whitelist.txt` file contents. Any updates to this file should be performed on all the stripes, as described in the following section.

Dynamic updates

After a cluster is started with whitelisting enabled, entries can be dynamically added to or removed from the whitelist without the need for server restarts. To perform a dynamic update, edit the `whitelist.txt` file contained in the server security root directories, and run the `ipwhitelist-reload` command to notify the servers in the cluster to reload the `whitelist.txt` file. Refer to the [“The “ipwhitelist-reload” Command” on page 78](#) section for more details.

Errors during whitelist reload, if any, are logged in respective server logs. Thus, after every update operation, server logs should be checked to verify that the updates took effect on all the servers.

If a cluster is not yet activated and the whitelist file needs to be reloaded on the servers, the server-level `ipwhitelist-reload` command can be used. It may also be helpful when the machine from where the cluster tool is to be used is itself not whitelisted initially. In this scenario, adding this machine's IP to the whitelist, and running the server-level `ipwhitelist-reload` command ensures that cluster tool can configure the cluster later.

Note:

If any failures happen while reading the `whitelist.txt` file during a dynamic update, the update is ignored and the server continues with the current whitelist. No partial updates are applied.

Connection Behaviour

When a client connects to a server, the server accepts the socket connection, and verifies the IP of the incoming client connection against the whitelist. If it finds that the client IP is not whitelisted, it closes the socket connection.

If a whitelisted client is removed from the whitelist via a dynamic update, it remains connected to the cluster as long as there is no network disconnection or explicit connection closure from the client. Subsequent connection attempts from the client to cluster will fail.

Full example of secure server configuration

Here is an example of a security configuration section in the config file of a single stripe, two node cluster, that configures auditing, encrypted connections, LDAP authentication and IP whitelisting:

```
authc=ldap
ssl-tls=true
whitelist=true
stripe.1.node.1.audit-log-dir=/path/to/audit-directory
stripe.1.node.1.security-dir=/path/to/security-root-directory
stripe.1.node.2.audit-log-dir=/path/to/audit-directory
stripe.1.node.2.security-dir=/path/to/security-root-directory
```

With the configuration in the above example, the server's security root directory would contain a truststore in the `trusted-authority` subdirectory, a keystore and a `credentials.properties` in the `identity` directory, an `ldap.properties` in the `access-control` subdirectory and a `whitelist.txt` file.

The `credentials.properties` file is required so that the server can connect to other servers in the stripe.

Client configuration

To configure a client to connect to a secured cluster, you need to give the client a path to the client's security root directory. This should contain, for example, the credentials that the client needs to connect to the cluster.

Command line tools

To enable command-line tools to connect to a secure cluster, a command must be prefixed with `-srd` (long option `--security-root-directory`).

The following example shows the use of the config tool `get` command with the `-srd` option specified:

```
./config-tool.sh -srd /path/to/security-root-directory get -c data-dirs -s localhost
stripe.1.node.1.data-dirs=main:%H/terracotta/user-data/main
```

Attempting to connect to a secure cluster without the `-srd` option will fail. Commands without this option retain their behavior.

Ehcache Client

An Ehcache client can define either an XML or a programmatic configuration, both of which support security configuration. The following are the examples of usages of each:

1. API Example

```
PersistentCacheManager cacheManager = CacheManagerBuilder
    .newCacheManagerBuilder()
    .with(EnterpriseClusteringServiceConfigurationBuilder
        .enterpriseSecureCluster(connectionURI,
            securityRootDirectoryPath) // 1
        .autoCreate())
    .build(true);
```

1	<p><code>EnterpriseClusteringServiceConfigurationBuilder enterpriseSecureCluster(URI, Path)</code> lets you create a <code>CacheManager</code> using a secure connection. The first argument is the URI of the Terracotta cluster, appended with the <code>CacheManager</code> name. The second argument is the Path to the client's security root directory.</p> <p><code>EnterpriseClusteringServiceConfigurationBuilder enterpriseSecureCluster(Iterable<InetSocketAddress>, String, Path)</code> serves the same function, with the added support for IPv6 addresses.</p>
---	---

Note:

If the URI or `Iterable<InetSocketAddress>` contains host names, make sure that they match the host names specified in the server certificates.

2. XML Example

```
<ehcache:config
  xmlns:ehcache="http://www.ehcache.org/v3"
  xmlns:tc="http://www.terracottatech.com/v3/terracotta/ehcache">
  <ehcache:service>
    <tc:cluster>
      <tc:connection url="{cluster-uri}/CM"
security-root-directory="{security-root-directory}"/>
      <tc:server-side-config auto-create="true"/>
    </tc:cluster>
  </ehcache:service>
</ehcache:config>
```

1	<p><code>security-root-directory</code> lets you specify the path to the client's security root directory. Not passing this option retains the behavior of communicating with an unsecured cluster.</p>
---	---

TC Store Client

1. API Example

```
DatasetManager datasetManager = DatasetManager.secureClustered(
  connectionURI, securityRootDirectoryPath) // 1
  .build();
```

1	<p><code>DatasetManager.secureClustered(URI, Path)</code> lets you create a <code>DatasetManager</code> using a secure connection. The first argument is the URI of the Terracotta cluster. The second argument is the Path to the security root directory which is to be used for the connection.</p> <p><code>DatasetManager.secureClustered(Iterable<InetSocketAddress>, Path)</code> serves the same function, with the added support for IPv6 addresses.</p>
---	---

Note:

If the URI or `Iterable<InetSocketAddress>` contains host names, make sure that they match the host names specified in the server certificates.

2. XML Example

```
<clustered xmlns=
  "http://www.terracottatech.com/v3/store/clustered"> <!--1-->
```

```

<cluster-connection>
  <server host="localhost" port="9410"/>
  <security-root-directory>/path/to/security-root-directory
  </security-root-directory>                                <!--2-->
</cluster-connection>
</clustered>

```

1	Declares a clustered DatasetManager configuration.
2	<code>security-root-directory</code> lets you specify the path to the client's security root directory. Not specifying this element retains the behavior of communicating with an unsecured cluster.

TMS Security

TMS security is configured using the `tmc.properties` file.

Setting the security root directory

If you need to set the security root directory to use for connections between a browser and the TMS, then you should set the `tms.security.root.directory` property to the path for the security root directory.

Windows platforms

Note that Windows paths often contain backslashes. The Java properties format requires backslashes to be escaped, so `C:\tcd\security-root-directory` would be configured using:

```
tms.security.root.directory=C:\\tcd\\security-root-directory
```

Alternatively, you can use forward slashes:

```
tms.security.root.directory=C:/tcd/security-root-directory
```

Encrypted connections between the browser and the TMS

You can enable HTTPS by setting:

```
tms.security.https.enabled=true
```

When HTTPS is enabled, you must also set the `tms.security.root.directory` property to the path for the security root directory. See [“Setting the security root directory” on page 119](#).

When HTTPS is configured, the specified security root directory must contain a valid truststore in the `trusted-authority` subdirectory and a valid keystore in the `identity` subdirectory.

When HTTPS is configured, the TMS is only accessible over HTTPS. Thus, any URL used to access the TMS must start with `https://`.

Browser warnings

If you are using self-signed certificates, you may see a security warning in your browser when connecting to the TMS. This is because the certificate was not signed by a certificate authority registered with the browser as a trusted certificate authority. We suggest that you check the certificate and verify its authenticity.

Furthermore, you should add the certificate to your browser's list of trusted root certificate authorities.

Note:

Each cookie entry is associated with a certain domain (not including port), and some browsers may remember the protocol under which the cookie has been set. So if you switch between http and https, or between localhost and 127.0.0.1, you should clear related cookies before reloading TMC in the browser.

Authentication

To configure TMS to require users to log in, you should set the `tms.security.authentication.scheme` property. There are two options: `file` and `ldap`, for file-based authentication and LDAP-based authentication respectively.

When authentication is configured, you must also set the following properties:

- `tms.security.root.directory` - see [“Setting the security root directory” on page 119](#).
- `tms.security.authorization.scheme` - see [“Authorization” on page 120](#).

File-based authentication requires a `users.xml` file to be added to the security root directory; LDAP-based authentication requires an `ldap.properties` file. See [“Authentication” on page 113](#) for more details.

Authorization

To control which access to TMS operations, you should set the `tms.security.authorization.scheme` property. There are three options: `authenticated`, `file` and `ldap`. The `authenticated` option allows access to all operations to all authenticated users. The `file` and `ldap` options correspond to file-based authorization and LDAP-based authorization respectively.

When authorization is configured, you must also set the following properties:

- `tms.security.root.directory` - see [“Setting the security root directory” on page 119](#).
- `tms.security.authentication.scheme` - see [“Authentication” on page 120](#).

TMS supports one role: `admin`. Users with no role can use most of the TMS functionality, but some operations, such as shutting down a server, require the `admin` role.

File-based authorization

File-based authorization requires a `users.xml` file to be added to the security root directory. See [“Authentication” on page 113](#) for more details.

For file-based authorization, the `users.xml` file must contain role information.

Example of a `users.xml` file for authorization:

```
<users>
  <user>
    <username>alex</username>
    <roles/>
  </user>
  <user>
    <username>beth</username>
    <roles>
      <role>admin</role>
    </roles>
  </user>
</users>
```

In this example `users.xml` file, `beth` has the `admin` role whereas `alex` does not.

Note that you can use a mixture of authentication and authorization schemes. For example, you could use LDAP-based authentication with file-based authorization. As in the example above, the `users.xml` file need not contain password hash information, if it is not being used for authentication. However, if you use `file` for both the authentication and authorization schemes, then you must specify both password hashes and role information.

Example of a `users.xml` file for both authentication and authorization:

```
<users>
  <user>
    <username>alex</username>
    <password>
      <algorithm>bcrypt</algorithm>
      <hash>$2a$10$UoM85/5I4SnIbr0QuFZ43ekffuQKSxZmL93bR9VMcdr2URmPyjyX2</hash>
    </password>
    <roles/>
  </user>
  <user>
    <username>beth</username>
    <password>
      <algorithm>bcrypt</algorithm>
      <hash>$2a$10$6D6c79lE0k/0SxrEtnfhGe2Yr.ygG0rFP1QzeyD9qshIMRrpUMOAS</hash>
    </password>
    <roles>
      <role>admin</role>
    </roles>
  </user>
</users>
```

LDAP-based authorization

LDAP-based authorization requires an `ldap.properties` file. See [“Authentication” on page 113](#) for more details.

The properties in the `ldap.properties` file are the same properties that are used to configure LDAP integration in other Software AG products. See [“LDAP Properties” on page 123](#) for a full list of supported properties.

Example of an `ldap.properties` file for authorization:

```
url=ldap://ldapserver.example.com:389
userrootdn=ou=People,dc=example,dc=com
uidprop=uid
personobjectclass=person
memberinfoingroups=true
mattr=uniqueMember
grouprootdn=ou=Group,dc=example,dc=com
gidprop=gid
groupobjectclass=group
```

You may wish to map the role defined in the LDAP server to the TMS `admin` role, in which case, you can add the `tcdb.roleMap` property to the `ldap.properties` file. For example:

```
tcdb.roleMap=terracottaTmsAdmin=admin
```

would map the `terracottaTmsAdmin` group defined in the LDAP server onto the TMS `admin` role.

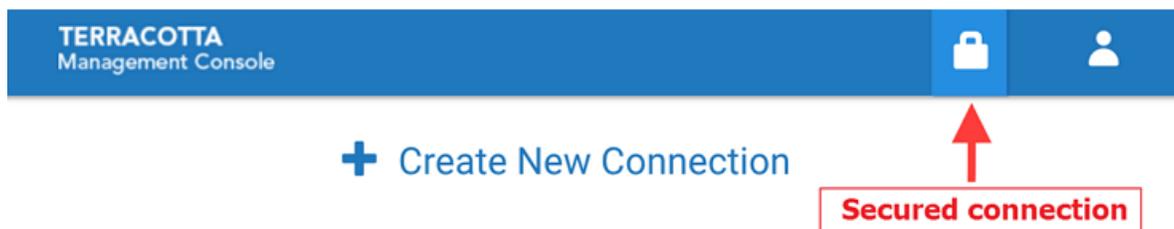
Auditing

To configure auditing for the TMS, set the `tms.security.audit.directory` property to the audit directory.

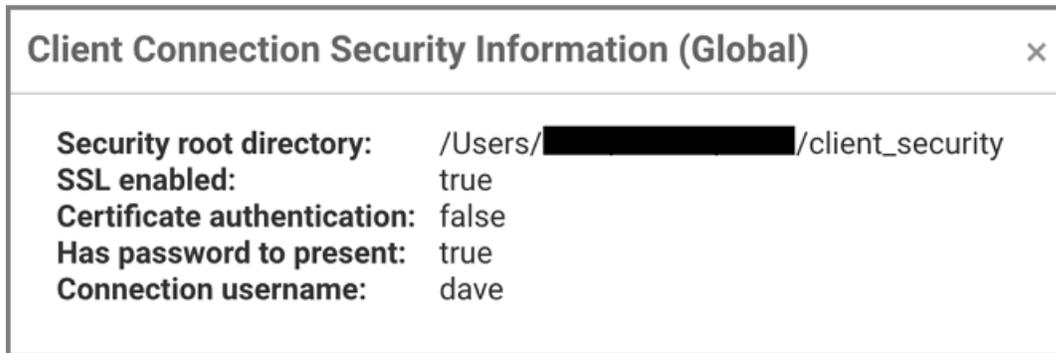
Connecting the TMS to secured clusters

To allow the TMS to connect to a secure cluster, you should set the `tms.security.root.directory.connection.default` property to the path for the security root directory containing, for example, credentials to connect to the secured cluster.

After configuring a secured cluster and setting up the TMS client security root directory, you should see a padlock icon in TMC:



When you click on the padlock, a pop-up window shows some security related information:

**Note:**

This security root directory and its configuration relates to the connection between the TMS server and the secured cluster. It is independent of the connection security between browser and TMS server.

SSL enabled	Connection is SSL enabled - i.e. the security root directory has a trusted authority certificate.
Certificate authentication	Connection is 2-way SSL enabled - i.e. the security root directory has an identity certificate.
Has password to present	The TMS/TMC client has a <code>credentials.properties</code> file.

Windows platforms

Note that Windows paths often contain backslashes. The Java properties format requires backslashes to be escaped, so `C:\tcd\audit-directory` would be configured using:

```
tms.security.audit.directory=C:\\tcd\\audit-directory
```

Alternatively, you can use forward slashes:

```
tms.security.audit.directory=C:/tcd/audit-directory
```

LDAP Properties

The minimum set of properties to specify is:

```
url  
userrootdn  
uidprop  
personobjclass
```

If you are using LDAP for authorization, you must also specify:

```
memberinfoingroups  
mattr
```

and, if you set `memberinfoingroups` to `true`, you must also specify:

```
grouprootdn  
gidprop  
groupobjclass
```

Most other properties need only be used if you have specific requirements.

Connection related properties

`url` - The URL of the LDAP server (e.g `ldap://ldapserver:389` or `ldaps://ldapserver:636`).

`keystoreUrl` - The URL from which a keystore can be retrieved (e.g. `file:///usr/local/ldap/keystore.jks`) - used to authenticate to the LDAP server.

`keystoreType` - The store type of the keystore (e.g. JKS).

`keystorePassword` - The password to verify the integrity of the keystore.

`keyAlias` - The alias in the keystore where the certificate and key are stored.

`keyPassword` - The password to allow access to the specified alias. Defaults to store password.

`truststoreUrl` - The URL from which a truststore can be retrieved. Used to to validate the certificate presented by the LDAP server during an SSL/TLS handshake.

`truststoreType` - The store type of the truststore (e.g. JKS).

`truststorePassword` - The password to verify the integrity of the truststore.

`noPrinIsAnonymous` - Set to true for LDAP servers that allow anonymous connections.

`prin` - The username to use to authenticate to the LDAP server.

`cred` - The password to use to authenticate to the LDAP server.

Note:

If `prin` and `cred` are not specified and `noPrinIsAnonymous` is not set to true, then the username and password of the user attempting to authenticate to the cluster / TMS will be used to authenticate to the LDAP server.

`watt.server.ldap.ignore.serverCertificateValidity` - If true, then invalid certificates presented by the LDAP server are ignored.

`watt.server.ldap.extendedProps` - Extra properties to add to the LDAP context. Format: `key1=value1;key2=value2`

`watt.server.ldap.retryCount` - How many times to retry a connection if it fails.

`watt.server.ldap.retryWait` - How many milliseconds to wait between connection retries.

Properties related to how to interact with the LDAP server

`timeout` - LDAP query timeout in milliseconds.

`watt.server.ldap.DNescapeChars` - A list of characters that should be escaped.

`watt.server.ldap.DNescapePairs` - A list of characters that should not be re-escaped.

`watt.server.ldap.DNstripQuotes` - If `false`, then quotes that get added when escaping are not striped from DNs.

`watt.server.ldap.DNescapeURL` - If `true`, then the start of a DN is escaped. This is useful for referrals when DNs can start with a URL.

`watt.server.jndi.searchresult.maxlimit` - The maximum number of results to return from an LDAP search. Zero means unlimited.

Properties related to the schema of a user

`userrootdn` - The DN under which users can be found (e.g. `ou=People,dc=example,dc=com`).

`uidprop` - The attribute on a user which contains the primary ID of the user (e.g. `uid`).

`personobjclass` - The LDAP schema class for users (e.g. `person`).

`useaf` - If `true`, then the `dnprefix` and `dnsuffix` properties should be used.

`dnprefix` - A string added to the beginning of a username for the LDAP lookup.

`dnsuffix` - A string added to the end of a username for the LDAP lookup.

Properties related to the schema of a group

`grouprootdn` - The DN under which groups can be found (e.g. `ou=Group,dc=example,dc=com`).

`gidprop` - The attribute on a group which contains the primary ID of the group (e.g. `gid`).

`groupobjclass` - The LDAP schema class for groups (e.g. `group`).

Properties related to how the schema connects users and groups

`group` - A role automatically given to every user.

`memberinfoingroups` - If `true`, then group membership is in the group definitions under the `grouprootdn`. If `false`, then group membership is in the user definitions under the `userrootdn`.

`matr` - The attribute on a user that specifies a group to which the user belongs OR the attribute on a group that specifies a user is a member. The semantics depends on the choice of `memberinfoingroups`.

`recursiveSearchDepth` - How deep to search for groups that are members of other groups.

Properties that Terracotta supports in addition to other Software AG products

`tcdb.roleMap` - A mapping from group names on the LDAP server to roles used in Terracotta. Format: `group1=tcdbRole1;group2=tcdbRole2`

Note:

Multiple LDAP groups can map to the same role.

SSL / TLS Troubleshooting guide

This document provides a list of the most commonly seen problems related to “Cluster Security” on [page 113](#), and their solutions:

Problem category: Host fails to start

This section describes the most commonly seen problems related to a host (server or a client) startup.

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:  
com.terracottatech.security.common.exception.SecurityConfigurationException:  
configured security-root-directory /path/to/security-root-directory  
does not exist
```

Diagnosis

The specified security root directory does not exist.

Action

Make sure that the directory exists and contains `identity` and `trusted-authority` directories with valid keystores and truststores in them.

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:  
com.terracottatech.security.common.exception.SecurityConfigurationException:  
identity directory doesn't exist in configured  
security-root-directory /path/to/security-root-directory
```

Diagnosis

The specified security root directory does not contain an `identity` directory.

Action

Make sure that the directory exists inside the security root directory and contains valid keystores.

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:  
com.terracottatech.security.common.exception.SecurityConfigurationException:  
trusted-authority directory doesn't exist in configured  
security-root-directory /path/to/security-root-directory
```

Diagnosis

The specified security root directory does not contain a trusted-authority directory.

Action

Make sure that the directory exists inside the security root directory and contains valid truststores.

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
No acceptable keystore files found in identity directory
/path/to/security-root-directory/identity
```

Diagnosis

Either of:

- identity directory does not contain any keystores.
- identity directory contains keystores, but their file names are not in the format `${common name}-${yyyyMMdTThmmss}.jks` (e.g. `com.organization.host-20180223T102319.jks`).

Action

Make sure that identity directory contains keystores which follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
No acceptable truststore files found in trusted-authority directory
/path/to/security-root-directory/trusted-authority
```

Diagnosis

Either of:

- trusted-authority directory does not contain any truststores.
- trusted-authority directory contains truststores, but their file names are not in the format `${common name}-${yyyyMMdTThmmss}.jks` (e.g. `trusted-authority-20180223T102319.jks`).

Action

Make sure that trusted-authority directory contains truststores which follow the truststore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:  
com.terracottatech.security.common.exception.SecurityConfigurationException:  
Tried to use the password terracotta_security_password to load the  
keystore file  
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks  
but that failed
```

Diagnosis

Latest keystore file does not have `terracotta_security_password` as its password, where *latest keystore file* is the keystore file with the latest timestamp string in the filename (e.g., `host-20180131T120830.jks` is considered newer than both `host-20170131T120830.jks` and `host-20180131T120822.jks`).

Action

Make sure the keystores follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:  
com.terracottatech.security.common.exception.SecurityConfigurationException:  
Tried to use the password terracotta_security_password to read the  
keystore entry with alias terracotta_security_alias in the keystore  
file  
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks  
but that failed
```

Diagnosis

Latest keystore file does not have `terracotta_security_password` as `terracotta_security_alias` entry password, where *latest keystore file* is the keystore file with the latest timestamp string in the filename (e.g., `host-20180131T120830.jks` is considered newer than both `host-20170131T120830.jks` and `host-20180131T120822.jks`).

Action

Make sure the keystores follow the follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:  
com.terracottatech.security.common.exception.SecurityConfigurationException:  
Unable to find required private key/certificate chain entry using  
alias terracotta_security_alias in keystore file  
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks
```

Diagnosis

Latest keystore file does not have `terracotta_security_alias` as certificate alias, where *latest keystore file* is the keystore file with the latest timestamp string in the filename (e.g.,

host-20180131T120830.jks is considered newer than both host-20170131T120830.jks and host-20180131T120822.jks).

Action

Make sure the keystores follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
Certificate in keystore file
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks
is expired
```

Diagnosis

Latest keystore file contains an expired certificate, where *latest keystore file* is the keystore file with the latest timestamp string in the filename (e.g., host-20180131T120830.jks is considered newer than both host-20170131T120830.jks and host-20180131T120822.jks).

Action

Make sure the keystores follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
Certificate in keystore file
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks
is not valid yet
```

Diagnosis

Latest keystore file contains a certificate with a future start date, where *latest keystore file* is the keystore file with the latest timestamp string in the filename (e.g., host-20180131T120830.jks is considered newer than both host-20170131T120830.jks and host-20180131T120822.jks).

Action

Make sure the keystores follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
The common name org.host of the certificate that was loaded from
```

```
keystore file
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks
doesn't match the common name com.organization.host in the filename
com.organization.host-20180131T120830.jks
```

Diagnosis

Common Name field in the Distinguished Name in the host's certificate does not match the common name fragment in the latest keystore filename.

Action

Make sure the keystores follow the keystore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
No valid trusted certificates found in trusted-authority directory
/path/to/security-root-directory/trusted-authority; Unable to find
required trusted certificate entry using alias
terracotta_security_alias in truststore file
/path/to/security-root-directory/trusted-authority/trusted-authority-20180131T120832.jks
```

Diagnosis

No valid truststores were found. The specific truststore reported in the Exception message contains a certificate which uses an alias other than `terracotta_security_alias`. Note that this Exception can be followed by one or more Suppressed Exceptions that can indicate why other truststores could not be used.

Action

Make sure the truststores follow the truststore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
No valid trusted certificates found in trusted-authority directory
/path/to/security-root-directory/trusted-authority; Certificate in
truststore file
/path/to/security-root-directory/trusted-authority/trusted-authority-20180131T120834.jks
is expired
```

Diagnosis

No valid truststores were found. The specific truststore reported in the Exception message contains an expired certificate. Note that this Exception can be followed by one or more Suppressed Exceptions that can indicate why other truststores could not be used.

Action

Make sure the truststores follow the truststore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
No valid trusted certificates found in trusted-authority directory
/path/to/security-root-directory/trusted-authority; Certificate in
truststore file
/path/to/security-root-directory/trusted-authority/trusted-authority-20180131T120834.jks
is not valid yet
```

Diagnosis

No valid truststores were found. The specific truststore reported in the Exception message contains a certificate with a future start date. Note that this Exception can be followed by one or more Suppressed Exceptions that can indicate why other truststores could not be used.

Action

Make sure the truststores follow the truststore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
No valid trusted certificates found in trusted-authority directory
security-root-client\trusted-authority; Tried to use the password
terracotta_security_password to load the truststore file
security-root-client\trusted-authority\trusted-authority-20180131T120832.jks
but that failed
```

Diagnosis

No valid truststores were found. The specific truststore reported in the Exception message does not have terracotta_security_password as its password. Note that this Exception can be followed by one or more Suppressed Exceptions that can indicate why other truststores could not be used.

Action

Make sure the truststores follow the truststore creation rules as described in the section [“Certificates” on page 109](#).

Symptom

Host fails to start with an Exception message similar to:

```
java.lang.RuntimeException:
com.terracottatech.security.common.exception.SecurityConfigurationException:
```

```
Unable to validate certificate chain with alias
terracotta_security_alias in keystore file
/path/to/security-root-directory/identity/com.organization.host-20180131T120830.jks
using truststore file(s)
```

Diagnosis

Host certificate in latest keystore file is not signed by any of the known trusted authorities, and thus cannot be validated by any of the truststore files. *Latest keystore file* here is the keystore file with the latest timestamp string in the filename (e.g., `host-20180131T120830.jks` is considered newer than both `host-20170131T120830.jks` and `host-20180131T120822.jks`).

Action

Make sure that the latest keystore file contained in the identity directory is signed by the truststores in the trusted-authority directory.

Problem category: Connection fails to establish

Symptom

```
org.terracotta.connection.ConnectionException:
com.terracottatech.connection.ProbableSecurityConfigurationException:
Handshake with server failed when this client tried to initiate a
non-secure connection. Possible reason: server is running with
security enabled.
```

Diagnosis

The client which tried to establish a connection to a server running with SSL/TLS configuration is not using an SSL/TLS configuration.

Action

Make sure the client uses a correct SSL/TLS configuration (via a secure API, an XML config, command parameters etc.) so that it can establish a secure connection with an SSL/TLS security-enabled server.

Symptom

```
org.terracotta.connection.ConnectionException:
com.terracottatech.connection.ProbableSecurityConfigurationException:
Handshake with server failed when this client tried to initiate a
secure connection. Possible reasons: client security configuration
is not valid, or server is not running with security
enabled.
```

Diagnosis

Either of:

1. The client is using an SSL/TLS security configuration but the server is not.
2. The client cannot validate the server because the server's CA certificate is not present in the client's trusted certificates.

3. The server cannot validate the client because the client's CA certificate is not present in the server's trusted certificates.

Action

Make sure that:

1. The client uses an unsecured configuration (via an unsecured API, an XML config, command parameters etc.) if the server is running with an unsecured configuration.
2. The client and the server certificates are signed by the same CA and their trusted-authority directories contain the same truststores.

26 Terracotta Server Migration from BigMemory to Terracotta

The new generation of Terracotta Server and platform has very significant changes from BigMemory with respect to handling of cluster topology, data storage formats, and various other aspects.

Because of this, there is no possibility of direct migration of data and configuration from a BigMemory installation to a Terracotta installation.

If you intend to reuse the same host systems in Terracotta that you were using in BigMemory, the Terracotta Server configuration file(s) from an existing BigMemory installation may be a useful reference when you create your new Terracotta configuration file(s). The BigMemory configuration file(s) contain the host names, addresses, etc. that you have been using so far.

27 Using Command Central to Manage Terracotta

Software AG Command Central is a tool that release managers, infrastructure engineers, system administrators, and operators can use to perform administrative tasks from a centralized location. It assists with configuration, management, and monitoring tasks in a simple and flexible manner.

Terracotta server instances can be managed from Command Central like other Software AG products. Both the Command Line and Web Interfaces of Command Central are supported.

Supported Commands

Terracotta supports the following Command Central CLI (Command Line Interface) commands:

1. Inventory

- `sagcc list inventory components` : Lists information about run-time components.
- `sagcc get inventory components` : Retrieves information about a specified run-time component.

2. Lifecycle

- `sagcc exec lifecycle` : Executes a lifecycle action against run-time components. See “[Lifecycle Actions for Terracotta](#)” on page 138 for Terracotta-specific information about Lifecycle Actions.

3. Monitoring

- `sagcc get monitoring state` : Retrieves the run-time status and run-time state of a run-time component.
- `sagcc get monitoring alerts` : Lists the alerts for a specified run-time component.
- `sagcc get monitoring runtimestatus` : Retrieves the run-time status of a run-time component.

4. Configuration

- `sagcc get configuration data` : Retrieves data for a specified configuration instance that belongs to a specified run-time component.

- `sagcc list configuration types`: Lists information about configuration types for the specified run-time component. See “[Supported Configuration Types](#)” on page 138 for Terracotta-specific information about configuration types.
- `sagcc list configuration instances`: Retrieves information about a specific configuration instance that belongs to a specified run-time component.

5. Diagnostics

- `sagcc list diagnostics logs`: Lists the log files that a specified run-time component supports.
- `sagcc get diagnostics logs`: Retrieves log entries from a log file. Log information includes the date, time, and description of events that occurred with a specified run-time component.

For information about Command Central CLI commands, see the Command Central Help.

Supported Configuration Types

Terracotta supports creating instances of the following configuration types:

- `JVM-OPTIONS`: The JVM memory settings for the Terracotta Server instance in `JAVA_OPTS` environment variable format.

Changes to this configuration will be effective upon a server restart.

- `TC-SERVER-NAME`: The name for the Terracotta Server instance. Editing this property is not allowed.

Lifecycle Actions for Terracotta

Terracotta supports the following lifecycle actions with the `sagcc exec lifecycle` CLI command and the Command Central Web Interface:

- `Start`: Start a server instance.
- `Restart`: Restart a running server instance.
- `Stop`: Stop a running server instance.

Runtime Monitoring Statuses for Terracotta

Terracotta can return the following statuses from `sagcc get monitoring runtimestatus` and `sagcc get monitoring state` CLI commands and the Command Central Web Interface:

- **Starting**: The server instance is starting. This is usually shown when:
 - The server was just started.
 - The server is a slave (Passive) synchronizing with its master (Active).
 - The server is recovering from an error condition.

- **Not Ready:** The server is not ready to accept client requests. To make it ready, follow the steps defined in the section [Making server ready](#).
- **Online Master:** The server instance is running and is the master (Active) in its stripe.
- **Online Slave:** The server instance is running and is a slave (Passive) in its stripe.
- **Stopping:** The server instance is stopping.
- **Stopped:** The server instance is not running.
- **Failed:** The server instance was running, but crashed or was killed without the SPM plugin's knowledge. If the server had crashed, checking its logs may help uncover the reason.
- **Unresponsive:** The server instance is running, but is not responding.
- **Unknown:** The state of the server instance is not known. This is most likely because of an unexpected exception or error that occurred while trying to fetch the server status.

Directory structure

The Terracotta server and SPM related files can be found under `${INSTALL_ROOT}/TerracottaDB/server/SPM`. This directory contains the following:

1. `bin`

Contains scripts to start and shut down the server.

2. `conf`

Contains the Terracotta config file `cluster.properties`. If any changes to the configuration are required, such as increasing the offheap, this file needs to be updated manually. See the [“Updating the config file” on page 140](#) section for more details.

Note:

This is the only directory in which content can be changed.

3. `instance`

Contains Terracotta SPM instance related metadata files.

4. `server-data`

Contains data maintained by Terracotta server. The contents of this directory are useful for troubleshooting problems with the server. The table below summarizes the contents of this directory:

Directory	Description
logs	Contains Terracotta server logs
metadata	Contains Terracotta server metadata
user-data	Contains Terracotta client data

Directory	Description
config	Contains Terracotta server configuration repository

Updating the config file

A Terracotta configuration file is a Java properties file containing configuration and topology information of the entire Terracotta cluster. Any changes to the config file must be done manually.

To find out more, visit [“The Terracotta Configuration File” on page 41](#).

Making configuration changes

Configuration changes in the default configuration `cluster.properties` file can be done freely before cluster activation. It's advised to make all the important changes before activating the cluster. After the cluster has been activated, the configuration can still be changed, but that's subject to certain constraints.

To find out more, visit [“Performing configuration changes” on page 51](#).

Making topology changes

Topology changes in the default configuration `cluster.properties` file can be done freely before cluster activation. If the configuration of a new node is added to this file, remember to use the same file across all installations. Also, the server name for the current installation must be updated in the `server-name` file (located next to `cluster.properties`).

The cluster topology can be changed dynamically as well.

To find out more, visit [“Adding nodes using the attach command” on page 49](#) and [“Removing nodes using the detach command” on page 50](#) in the description of the config tool.

Making the server ready

When the Terracotta server is started for the first time from Command Central, it will be in the 'Not Ready' state. To make the server ready to accept client requests, it needs to be part of an activated cluster. The `activate` command of the config tool (located under `${INSTALL_ROOT}/TerracottaDB/tools/bin`) can be used to activate the cluster.

See the config tool section [“Activating a cluster using activate command” on page 48](#) for more details.

28 Terracotta in Network Environments with Subnets

If Terracotta nodes reside in a given subnet (for example in a Kubernetes cluster) and clients in another (for example outside of the Kubernetes cluster), node addresses (host name or IP address) are not resolvable by clients. Administrators of the Terracotta cluster can use public addresses in this scenario to assign public names to Terracotta cluster nodes, using which clients can establish connections.

Configuring public addresses

Public addresses, if configured, need to be defined on all the nodes in the cluster. This can be achieved using one of the following ways:

Using command-line parameters during node startup

Public addresses can be configured on each node during startup by using the properties `public-hostname` and `public-port`.

The following example illustrates setting the public hostname `tc-cluster.public.com` and public ports 1111 and 2222 respectively on two nodes during startup:

```
> start-tc-server.sh public-hostname=tc-cluster.public.com public-port=1111 -n node-1
...
> start-tc-server.sh public-hostname=tc-cluster.public.com public-port=2222 -n node-2
...
```

Using config file during node startup

Public addresses for the entire cluster can alternatively be saved in a Terracotta config file, which can later be used to start a cluster. See the section [“The Terracotta Configuration File” on page 41](#) for more details.

The following example illustrates setting the public hostname `tc-cluster.public.com` for both the stripes, and public ports 1111 and 2222 for the first and second stripes respectively:

```
stripe.1.node.1.public-hostname=tc-cluster.public.com
stripe.1.node.1.public-port=1111
stripe.2.node.1.public-hostname=tc-cluster.public.com
stripe.2.node.1.public-port=2222
```

Using config tool "set" command

The config tool set command can be used to dynamically configure public addresses on a cluster. NOTE: Public addresses can be updated without restarting the cluster. See the section [“Performing configuration changes” on page 51](#) for more details.

The following example illustrates setting the public hostname `tc-cluster.public.com` on all the nodes of a 2 stripe cluster, and public ports 1111 and 2222 respectively:

```
> config-tool.sh set -s tc-cluster.internal.com:9410
-c public-hostname=tc-cluster.public.com
-c stripe.1.node.1.public-port=1111
-c stripe.2.node.1.public-port=2222
```

SSL/TLS considerations The SSL/TLS certificates of the Terracotta nodes will need to include Subject Alternative Names (SANs) that match the public addresses.

Using public addresses

Once the public addresses have been set, all tools (e.g. config tool, cluster tool etc.) would use the public addresses by default.

The following example illustrates the execution of the cluster tool status command on the cluster with public addresses configured. Take note of how the internal to public address mapping is displayed in the Address column:

```
> cluster-tool.sh status -n tc-cluster -s tc-cluster.public.com:1111
| STRIPE: 1 |
+-----+
| Server Name | Address |
Status |
+-----+
| node-1 | tc-cluster.internal.com:9410=tc-cluster.public.com:1111 |
ACTIVE |
+-----+
| STRIPE: 2 |
+-----+
| Server Name | Address |
Status |
+-----+
| node-2 | tc-cluster.internal.com:9510=tc-cluster.public.com:2222 |
ACTIVE |
+-----+
```