

HP Operations Orchestration

For the Windows® and Linux operating systems

Software Version: 9.02

High Availability Guide

Document Release Date: December 2011

Software Release Date: December 2011



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2008- 2011 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at:

<http://www.hp.com/go/hpsoftwaresupport>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

High Availability Guide.....	1
Contents.....	5
Overview.....	7
Best Practices for Setting up High Availability in HP Operations Orchestration.....	7
Introduction to High Availability in HP Operations Orchestration.....	7
Clustering Service Overview.....	7
Horizontal Scalability of Central Servers.....	8
Planning High Availability.....	9
OO Cluster Behavior in Node Failures.....	10
Recommended Deployment for the HP Operations Orchestration High Availability Solution.....	12
Perform a Clean Installation.....	14
Converting a Standalone Central to a Clustered Central.....	16
Add a Clustering Service to an Existing Central on Windows.....	16
Add a Clustering Service to an Existing Central on Linux.....	17
Adding a Node to an Existing Cluster.....	21
Adding a Node to an Existing Cluster Overview.....	21
Installing Additional Central Instances in Windows.....	22
Installing Additional Central Instances in Linux.....	23
Load Balancer Configuration.....	26
Providing Load Balancing for OO Components.....	26
Providing High Availability to the Load Balancer.....	27
Installing the HP OO Load Balancer on a Windows Platform.....	27
Installing the HP OO Load Balancer on a Linux Platform.....	28
Configuring a Load Balancing Cluster.....	28
Connecting Studio to the Load Balancer.....	32
Upgrading the Load Balancer.....	32
Scheduler Configuration.....	33

Introduction	33
Working with the Scheduler in a Cluster	33
Create a Clustered Scheduler	33
Configure the Scheduler to Work With a Load Balancer	34
Common Tasks	35
Prepare the Shared File System	35
Verify That the Central Cluster is Correctly Configured	36
Advanced Clustering Configuration	39
Removing a Node from an OO Cluster	39
Remove a Node from a Windows Cluster Installation	39
Remove a Node from a Linux Cluster Installation	40
Install a Standalone Clustering Service	41
Troubleshooting	42
Split Brain Problem	42
Clustering Service Problems	45
Terracotta Dirty Database Problem	46
NFS Problem	46
Corrupted Repository Problem	47

Chapter 1

Overview

This chapter includes:

Best Practices for Setting up High Availability in HP Operations Orchestration	7
Introduction to High Availability in HP Operations Orchestration	7
Clustering Service Overview	7
Horizontal Scalability of Central Servers	8
Planning High Availability	9
OO Cluster Behavior in Node Failures	10

Best Practices for Setting up High Availability in HP Operations Orchestration

Note: This guide supports HP Operations Orchestration 9.00.01 and higher.

This guide contains a general overview of the High Availability solution, and describes the process of setting up your High Availability environment in Windows and Linux from different possible starting points.

You can:

- Deploy a fresh HP Operations Orchestration clustered environment. For details, see ["Perform a Clean Installation" \(on page 14\)](#).
- Expand a standalone HP Operations Orchestration Central to a clustered environment. For details, see ["Converting a Standalone Central to a Clustered Central" \(on page 16\)](#).
- Scale out a HP Operations Orchestration clustered environment to increase throughput capacity and high availability. For details, see ["Adding a Node to an Existing Cluster" \(on page 21\)](#).

Introduction to High Availability in HP Operations Orchestration

If you run large volumes of flows, and if high availability is a requirement for your organization:

- An OO Central cluster can provide horizontal scalability, high availability of Central servers via failover protection, and automatic recovery of failed flows when a Central node fails.
- HP OO Load Balancer provides load balancing for requests to Central and RAS installations.

Clustering Service Overview

HP Operations Orchestration's clustering is based on a Clustering Service that controls all aspects related to clustered state and coordination. The Clustering Service is a separate entity from a

clustered Central process. It is a server process that can co-exist with a Central process on the same node, or it can be hosted on a different machine.

The clustering support is based on a JVM level clustering solution called Terracotta (for details, see www.terracotta.org). The OO Clustering Service is a Terracotta server in either active or passive mode.

Note: Central processes themselves use active/active topology. Only the Clustering Services use an active/passive topology.

All that is required for the proper functioning of an HP OO cluster is the existence of the active Clustering Service. The additional passive service is recommended for achieving high availability in case the active service fails. Any clustered transaction must be acknowledged by all standby services before the active Clustering Service acknowledges it to the involved Central services. This means that a large number of passive Clustering Services will decrease throughput. For most deployments, one active and one passive Clustering Service are sufficient.

Tip: It is recommended that both the active and the passive Clustering Services run on identical hardware.

When the cluster starts, the Clustering Services go through an election process, and one of them becomes the active instance.

Horizontal Scalability of Central Servers

When you install the Clustering Service and create a cluster of OO Central installations, the clustering technology has the following features:

- Central cluster uses an active-active topology, thus maximizing the utilization of available hardware.

Note: Clustering Services use active-passive topology.

- Horizontal scalability. The cluster can be horizontally scaled by adding more nodes. Routing all Studio and Central UI (browser) sessions, as well as headless requests through a load balancer is the recommended approach.
- The cluster uses one shared repository, and there is no need to manually synchronize the repository for each node. The shared Central repository requires a shared file system such as a Windows share (for Windows deployments) or a Network File System (NFS) share (for Linux deployments). If high availability of the repository is required, this shared file system should be implemented on top of your choice of HA file system technologies (for instance, SAN).

Caution: Do not start Central services simultaneously. If you have more than one Central service to restart, first restart the first one, and make sure it is up by logging in to its Web UI directly. Then, restart the second one, and so on.

Starting several Central services simultaneously might cause corruption in the shared repository. In addition, it might cause a problem when using NFS repository, as described in ["NFS Problem" \(on page 46\)](#).

Planning High Availability

When you are planning your high availability solution, you should consider the following:

- Should I cluster my Clustering Service installation for high availability?

The deployment of a passive Terracotta Service is not strictly required but is strongly recommended, since if you do not provide high availability for the Clustering Service, you have simply moved your vulnerability to a single point of failure from the Central installation to the Clustering Service installation. The passive (standby) Terracotta service provides high availability in case of the active Terracotta service's failure.

Be cautious, however, of deploying more passive Terracotta installations than are necessary for high availability. Each additional passive Terracotta installation comes with an operational cost. In a cluster, any clustered transaction must be acknowledged by all standby services before the active Clustering Service acknowledges it to the involved Central services. This means that each additional passive Clustering Service you install decreases throughput. For most deployments, one active and one passive (standby) Clustering Service are sufficient.

Note: It is recommended that both the active and the passive (standby) Clustering Service instances run on identical hardware.

- Should I load balance my OO environment?

For Central and standalone RAS servers, the OO Load Balancer provides Studio and Central UI balancing of sessions as well as balancing of headless flow invocations. The OO Load Balancer is based on the Apache Load Balancer. For more information on the Apache Load Balancer, see the Apache Software Foundation web site.

Note: To achieve high availability, you must use a load -balancing solution between the Centrals. You have the option of using a load balancing solution other than the OO Load Balancer.

- Should I have high availability failover protection of the Central database?

The Central database contains the Central nodes' run histories, schedules, and system configuration. If the database server fails without failover protection, Central and its users lose all of that information.

- Should I provide high availability failover protection of the load-balancer?

It is recommended that you provide high availability for your load-balancing solution. If you do not, this component becomes a single point of failure.

When the Load Balancer and OO Clustering Service are installed in a development environment, multiple authors developing flows have the benefits both of load balancing for their requests to the Central cluster and failover protection for each author in case the Central server node to which you are connected fails.

Note: When planning an OO deployment that avoids single points of failure, notice the underlying hardware as well. For example, if the OO servers are Virtual Machines, they should not be hosted on a single ESX.

Given the preceding recommendations, the "[Recommended Deployment for the HP Operations Orchestration High Availability Solution](#)" (on page 12) is the recommended configuration for your OO deployment. See also the following sections for more considerations with which to design high availability and load-balancing into your OO environment.

OO Cluster Behavior in Node Failures

There are a few types of node failures that affect the cluster behavior:

- **The Central service.** If a node in the Central cluster fails:
 - The other Central continues functioning, and when the failed Central instance comes back up, it joins the cluster and is ready for work.
 - The load balancer identifies that the node is down, and no longer directs flow executions to the Central service.
 - The failed Central node's flow executions are automatically recovered on another node of the cluster and resumed at their last check-pointed step.

Note: This feature is not enabled by default. It can be enabled by an OO administrator in the **Administration > System Configuration** tab. Only unassisted (headless) flow runs can be recovered by this feature. For details see HP Operations Orchestration Central User Guide.

- Any Central UI session connected to the failed node is closed and a new one has to be opened (if Central UI sessions always go through a load balancer (recommended), then only logging on again is required). Flows that were initiated through that session can be manually resumed from the Central Administration tab once a session to another node is established.
- Any Studio connection to the Central public repository that had gone through the failed Central node is closed. Ongoing editing work is not lost, as it is stored in the shared repository, in a private workspace.

Note: The system maintains edited work that has been saved. When the connection to a Central node is lost, edited work done in Studio is not recovered if it was not saved.

- **The Clustering Service.** If the Clustering Service fails on one node:
 - If the failed Clustering Service was the active node, the passive (standby) node takes over and becomes active. The Central instances resume functionality after this failover is completed.
 - If the failed Clustering Service was the passive node, the cluster continues functioning, without a passive standby. Further failure of the active instance results in cluster "death" and requires the restart of all nodes.

- **The Load Balancer.** Once the load balancer is down, the various OO cluster components cannot be accessed through it. In such a case, the OO cluster components can only be accessed directly.

Chapter 2

Recommended Deployment for the HP Operations Orchestration High Availability Solution

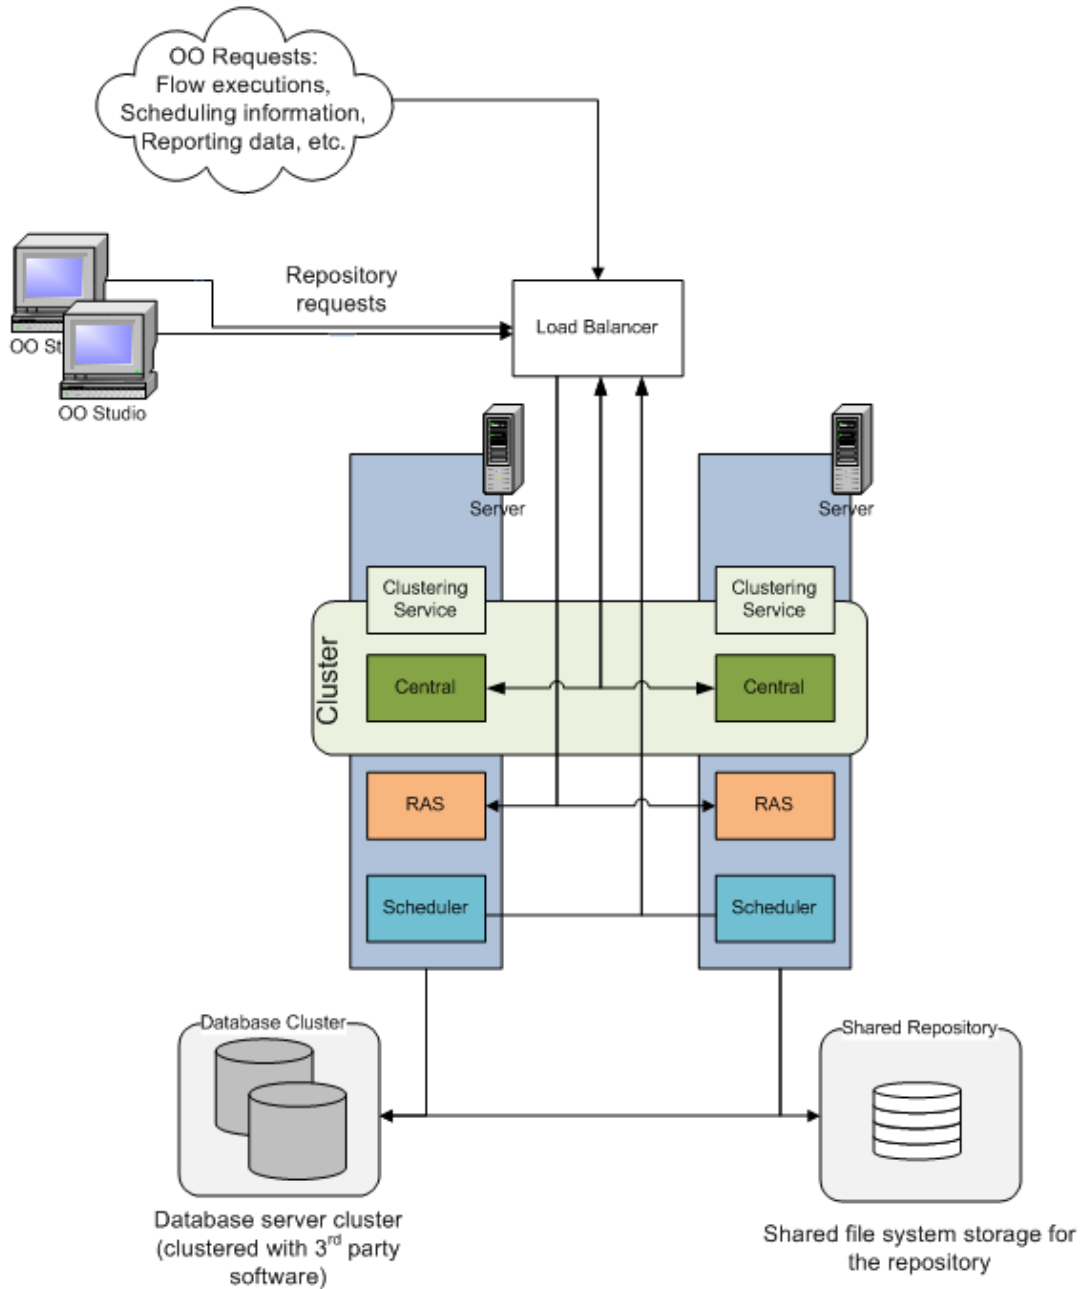
The recommended deployment configuration comprises the following:

- A two-node Central cluster, in which each node hosts one Central instance and one Clustering Service instance, on identical hardware
- Highly available shared file system, for hosting the shared repository.
- A two-node database cluster, for hosting the Central database
- A load-balancing solution, for load-balancing requests to Central and RAS instances

The diagram below describes the recommended configuration for a high availability solution.

High Availability Guide

Chapter 2: Recommended Deployment for the HP Operations Orchestration High Availability Solution



Chapter 3

Perform a Clean Installation

To create a Central cluster with the OO Load Balancer, perform the following steps:

1. Create the database cluster. How you create the database cluster depends on the third-party software that you use to create and support it.
2. Create a shared file system. For details, see ["Prepare the Shared File System" \(on page 35\)](#).

Tip: It is recommended not to use any of the cluster nodes for hosting the shared file system.

3. Install two standalone Centrals on the desired nodes that point to the same database. This is done as follows:
 - a. Install the first Central (using a clean database or pointing to an existing one). For details, see the Installing or Upgrading HP Operations Orchestration Guide.
 - b. Install the second Central on the other node. During the installation, choose the option of using the database defined in the previous step.

Note:

- Make sure that every Central instance is functional.
- If you clustered your database servers with an Oracle Real Application Cluster, then be sure that you have installed OO Central according to the provisions described in the Installing Operations Orchestration with Oracle Real Application Clusters Guide).
- After installing a standalone Central, make sure your Central is upgraded (to 9.02)
- When installing the central on each node, you need to select the **Create or recreate (if it exists), the database schema** box only once.

4. Give Read/Write file-system level access to Central processes to the share. This step depends on the chosen type of the shared file system.
5. Create a Clustering Service on each Central node. For Windows, see ["Add a Clustering Service to an Existing Central on Windows" \(on page 16\)](#) . For Linux, see ["Add a Clustering Service to an Existing Central on Linux" \(on page 17\)](#).
6. Restart each Central Service, one at a time.
7. Install the OO Load Balancer. For information on installing the OO Load Balancer by itself, see ["Load Balancer Configuration" \(on page 26\)](#).
8. Configure the Scheduler in the Centrals. For details, see ["Working with the Scheduler in a Cluster" \(on page 33\)](#).
9. Integrate a load balancer with node management from Central. For details, see the section

"Integrating a Load balancer with Node Management" in the HP Operations Orchestration Central User Guide.

Note: Modify the load balancer URL in Studio by adding the `lbUrl` as follows:
`https://<LB_ip_address>:<LB_port>/balancer-manager.`

Chapter 4

Converting a Standalone Central to a Clustered Central

To convert an existing standalone Central to Clustered Central:

1. Create a new shared folder for the Shared repository. For details, see "[Prepare the Shared File System](#)" (on page 35).
2. On the second node, install a Central (using the existing database and shared repository). For details, see the Installing or Upgrading HP Operations Orchestration Guide.
3. **For Windows only.** Configure the Central service to log on with credentials of a user which can access the Shared Repository.
4. On both nodes, install a Clustering Service and configure OO Central in clustering mode.
 - For Windows, see "[Add a Clustering Service to an Existing Central on Windows](#)" (on page 16)
 - For Linux, see "[Add a Clustering Service to an Existing Central on Linux](#)" (on page 17)
5. Restart each Central Service, one at a time.
6. **Important!** Verify that the cluster of Centrals is correctly configured. For details, see "[Verify That the Central Cluster is Correctly Configured](#)" (on page 36).
7. Configure the Load Balancer.
 - For Windows, see "[Installing the HP OO Load Balancer on a Windows Platform](#)" (on page 27).
 - For Linux, see "[Installing the HP OO Load Balancer on a Linux Platform](#)" (on page 28).
8. Configure the Scheduler in the Centrals. For details, see "[Working with the Scheduler in a Cluster](#)" (on page 33).

Add a Clustering Service to an Existing Central on Windows

Note: The installation path for the Clustering Services must not contain non-English language characters.

1. On the HP OO Installation CD, navigate to and double-click the **ClusterInstaller-9.00.exe** file.
2. On the **Welcome** page of the Setup wizard, click **Next**.
3. On the **License Agreement** page, accept the terms of the license agreement, and click **Next**.
The **Installation Type** page appears, presenting the various types of installation and configuration options that you will use to create a Central cluster.
4. Select **Install Clustering Service AND Configure OO Central in Clustering Mode**, and

then click **Next**.

The **Shared Repository Path** page appears.

5. In the **Shared Repository Path** page's text box, type the location of the shared folder, and click **Next**.

Or

Click **Browse**, navigate to the shared folder, and click **Next**.

For example, `\\my-server\shared_repo`.

The **Clustering Service High Availability Information** page appears.

6. To specify additional Clustering Service nodes (that is, the other machines that will host Clustering Services), select the **Configure Additional Clustering Services for High Availability** box.
7. In the **List of comma-delimited Clustering Service machine addresses** box, type a comma-delimited list of the machine names of the other Clustering Service nodes in the cluster, including those that will be part of the cluster but that have not had the Clustering Service installed yet.

For example, `abvc01.dev.com,abvc02.dev.com,abvc10.dev.com`.

Note:

- When typing the machine names, be sure that you use the machine names rather than the IP addresses.
- Make sure you type only the names of the machines that host Clustering Services.
- Make sure that case sensitivity is identical on all nodes in the cluster.

8. Click **Next**. The **Ready to Install** page appears.
9. On the **Ready to Install** page, click **Install**.

Note: At this stage, the contents of the current repository of the standalone Central are copied to the shared repository.

10. When the installation has finished, on the **Completing** page, click **Finish**.
11. Open the `<OO installation directory>\Clustering\terraccotta\wrapper.conf` and add the following line in the **Java Additional Parameters** section:

```
wrapper.java.additional.9=-Dcom.tc.l2.nha.dirtydb.autoDelete=true
```

Add a Clustering Service to an Existing Central on Linux

Note: The installation path for the Clustering Services must not contain non-English language characters.

To install the Clustering Service and configure the Central node to run in cluster mode:

1. On the installation CD, in the Clustering subdirectory, locate and copy the Cluster-9.00.zip file and extract it to the machine on which you're going to install the Clustering Service.

Extracting the .zip file creates the **Cluster-9.00** directory.

Note: The following steps create the information in **INSTALL.txt** (contained in the **Cluster-9.00.zip** file).

2. In the **Cluster-9.00** directory, open the **install.config** file for editing and modify it to specify the following.

To specify this	In this line of install.config	Type this between the quotation marks (“”)
Central home directory	ICONCLUDE_HOME=""	The location of the HP OO home directory (the location to which Central-9.00.zip was extracted).
Clustering Service home directory	ICONCLUDE_CLUSTER_HOME=""	The location of the Clustering Service home directory (the location to which Cluster-9.00.zip was extracted).
Installation of the Clustering Service and configuration of the server as a Clustering Service node	INSTALL_TYPE=""	terracotta_server
Clustering Service server	TS_SERVER=""	The machine name (on which you are installing the Clustering Service). The fully qualified machine name should be used, and should not be localhost.
The location of the shared Central repository	SHARED_REPO_PATH=""	Location of the shared Central repository, as referenced by the mount point. This location is the shared directory on the shared storage device.
Any additional Clustering Service servers	ADDITIONAL_TERRACOTTA_SERVERS_LIST=""	Comma-separated list of additional machine names hosting Clustering Service instances (for Clustering Service high availability).

For example, suppose that:

- You are installing Clustering Service on and configuring vancouver.mydomain.com as a Clustering Service node (TS_SERVER="vancouver.mydomain.com")
- You are clustering the Clustering Service with two other cluster nodes, tacoma.mydomain.com and portland.mydomain.com.

When you have configured the **install.config** file properly on `vancouver.mydomain.com`, it will look like the following:

```
ICONCLUDE_HOME="Central-9.00"

# the location where OO Cluster is installed. Make sure that you points to the C
lustering's home folder
ICONCLUDE_CLUSTER_HOME="Cluster-9.00"

# Installation types:
# "terracotta_central" --> Configure OO node in Clustering Mode
# "terracotta_server" --> Install Clustering Service
# "load_balancer" --> Install Cluster Load Balancer

INSTALL_TYPE="terracotta_server"

#Terracotta server host name or IP address (do not use localhost, 127.0.0.1, etc
)
TS_SERVER="vancouver.mydomain.com"

# The path for the clustered Repository
SHARED_REPO_PATH="/opt/sharedRepo"

# List of comma-delimited Clustering Service machine addresses if exist
ADDITIONAL_TERRACOTTA_SERVERS_LIST="tacoma.mydomain.com,portland.mydomain.com"
```

3. Run the **configure.sh** script.

The **configure.sh** script creates **CLUSTER.sh** in the `/bin` folder.

4. Edit the **install.config** file as follows: **INSTALL_TYPE="terracotta_central"**.
5. Run the **configure.sh** script again.

At this stage, the contents of the current repository of the standalone Central is copied to the shared repository.

Note: For each Clustering Service node that you create:

- The machine name in the `TS_SERVER=""` line will be that machine's name.
- The list of machines in the `ADDITIONAL_TERRACOTTA_SERVERS_LIST=""` line will differ, each time including only the nodes other than the one on which you are working.

For instance, when you configure **install.config** on `vancouver.mydomain.com`:

- The line `TS_SERVER=""` should specify `"vancouver.mydomain.com"`
- The list in the `ADDITIONAL_TERRACOTTA_SERVERS_LIST=""` line should be: `"tacoma.mydomain.com,portland.mydomain.com"`

6. To start Central, run the following command:

```
./bin/Central.sh start
```

7. To start the Clustering Service, run the following command:

```
./bin/CLUSTER.sh start
```

8. Repeat this procedure on each of the other nodes hosting Clustering Services.
9. Edit the **install.config** file as follows: **INSTALL_TYPE="terracotta_central"**.

Chapter 5

Adding a Node to an Existing Cluster

This chapter includes:

Adding a Node to an Existing Cluster Overview	21
Installing Additional Central Instances in Windows	22
Installing Additional Central Instances in Linux	23

Adding a Node to an Existing Cluster Overview

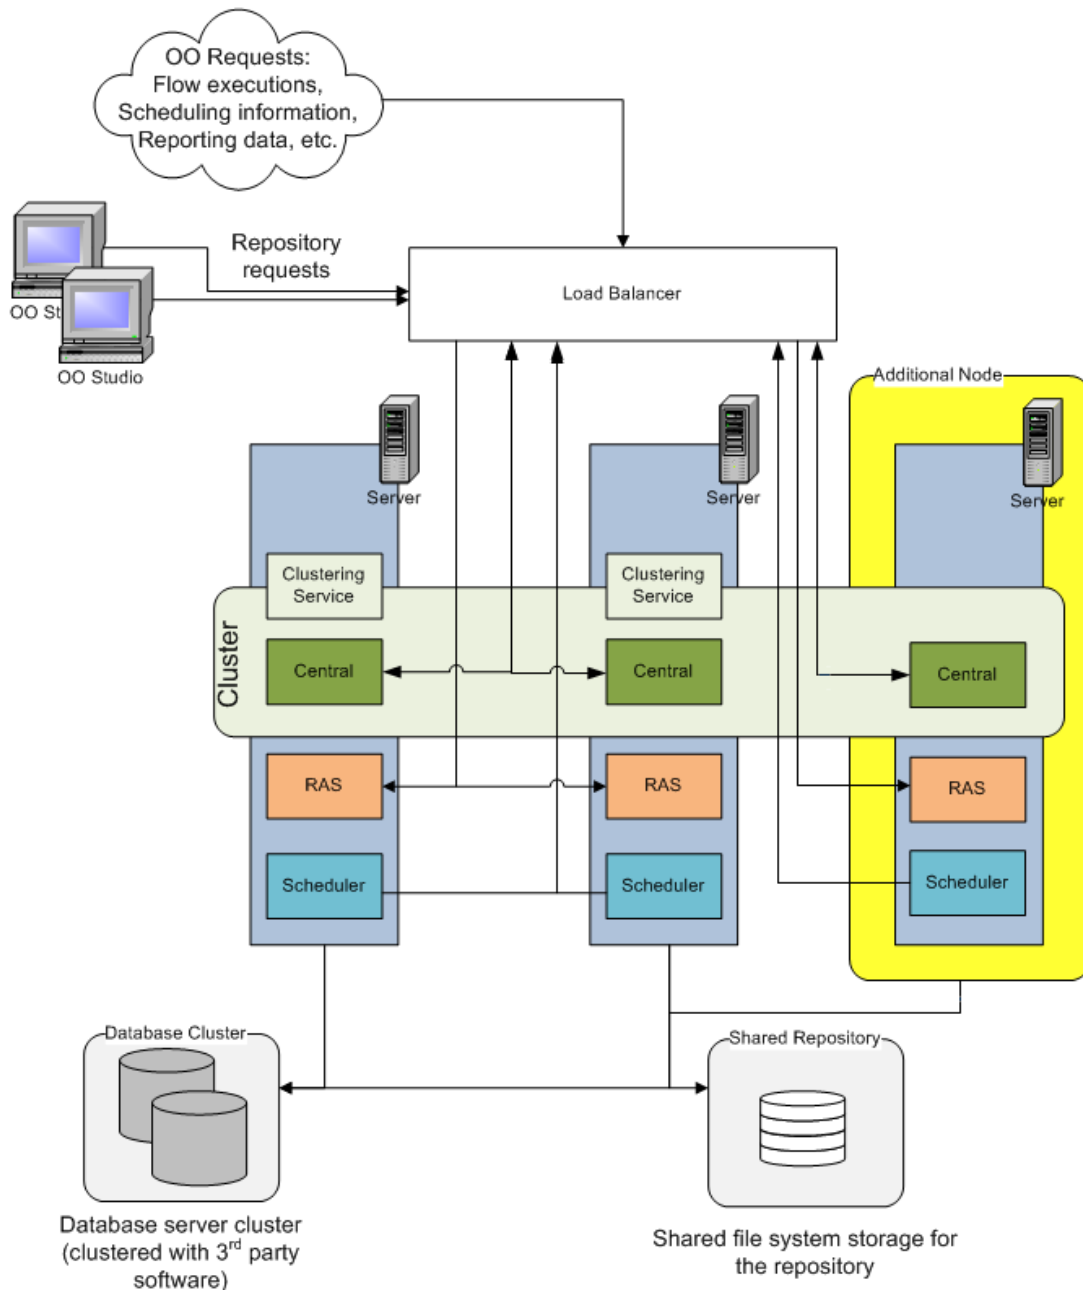
The best practices deployment for the HP Operations Orchestration HA solution (for details, see ["Recommended Deployment for the HP Operations Orchestration High Availability Solution" \(on page 12\)](#) achieves high availability from the perspective of Central and the Clustering Service. Also, because clustered Central instances always work in an active-active mode, it maximizes the usage of the hardware necessary to achieve HA.

When two Central nodes cannot cover the capacity requirements, you scale out the cluster by adding additional Central instances. This is achieved by installing additional Centrals and then switching them to clustering mode (see the installation procedures in subsequent sections). If more redundancy is desired for the Clustering Services themselves, some of the added nodes can also host Clustering Service instances (one per node). A large number of standby Clustering Service instances would decrease the throughput, so in general, we recommend that the added nodes just host clustered Central instances.

Caution: Multiple clustered Central processes running on the same node is not supported in all platforms. While in Linux it is technically possible to run multiple Central instances on the same node (physical machine or VM), if those Central processes are also clustered, the run recovery will not function properly due to the fact that they will all share the same node ID.

To add a node to an existing cluster, you install a Central, configure it to run in cluster mode, and connect the load balancer to it. For details, see ["Installing Additional Central Instances in Windows" \(on page 22\)](#) and ["Installing Additional Central Instances in Linux" \(on page 23\)](#).

The following diagram shows a three-node Central cluster with two Clustering Services.



Installing Additional Central Instances in Windows

This task describes how to add additional Central instances to the cluster without installing a Clustering service on it.

1. Install Central on the node to be added to the cluster. For details, see the Installing or Upgrading HP Operations Orchestration Guide.

Note: Verify that the installed Central has the same version as the other Centrals in the cluster.

2. Navigate to and double-click the file **ClusterInstaller-9.00.exe**.

3. In the **Welcome** page of the Setup wizard, click **Next**.
4. In the **License Agreement** page, accept the terms of the license agreement, and click **Next**. The **Installation Type** page appears, presenting the various types of installation and configuration options that you will use to create a Central cluster.
5. Select **Configure OO Central in Clustering Mode**, and then click **Next**. The **Shared Repository Path** page appears.
6. In the text box, type the location of the shared folder, and click **Next**.

Or

Click **Browse**, navigate to the shared folder, click **OK**, and then click **Next**.

The **Clustering Service Settings** page appears.

7. In the text box, type the comma-delimited list of machine hostnames that are hosting Clustering Services, and then click **Next**.

Note:

- When typing the machine names, be sure that you use the machine names rather than the IP addresses.
- Make sure you type only the names of the machines that host Clustering Services.
- Make sure that case sensitivity is identical on all nodes in the cluster.

8. On the **Ready to Install** page, click **Install**.
9. When the installation has completed, on the **Completing** page, click **Finish**.
10. **Important!** Verify that the Central cluster and Clustering Service cluster are correctly configured. For details, see "[Verify That the Central Cluster is Correctly Configured](#)" (on page 36)
11. Verify that the **RSCentral** service is up. If it is down, restart the service.
12. Connect the load balancer to the Central. For details, see "[Configuring a Load Balancing Cluster](#)" (on page 28).
13. Configure the Scheduler in the Central. For details, see "[Working with the Scheduler in a Cluster](#)" (on page 33).

Installing Additional Central Instances in Linux

This task describes how to add additional Central instances to the cluster without installing a Clustering service on it.

1. Install Central on the node to be added to the cluster. For details, see the Installing or Upgrading HP Operations Orchestration Guide.
2. Configure the shared file system. For details, see "[Prepare the Shared File System](#)" (on page 35).
3. On the installation CD, in the Clustering subdirectory, locate and copy the **Cluster-9.00.zip** file and extract it to the machine that hosts the Central Service that is being configured in

clustering mode.

Extracting the .zip file creates the Cluster-9.00 directory.

Note: The following steps create the information in **INSTALL.txt** (contained in the **Cluster-9.00.zip** file).

- In the Cluster-9.00 directory, open the **install.config** file for editing and modify it to specify the following.

To specify this	In this line of install.config	Type this between the quotation marks (“”)
OO Central home directory	ICONCLUDE_ HOME=""	The location of the Central home directory.
Clustering Service home directory	ICONCLUDE_ CLUSTER_ HOME=""	The location of the Clustering Service home directory (the location OS where Cluster-9.00.zip was extracted).
Configuration of the Central server as a Central node	INSTALL_ TYPE=""	terracotta_central
Clustering Service server	TS_SERVER=""	Machine name of Clustering Service host server (first Clustering Service node if more are configured for HS).
Repository shared by Central nodes	SHARED_REPO_ PATH=""	Location of the shared Central repository, as referenced by the mount point. This location is the shared directory on the shared storage device.
Any additional Clustering Service servers	ADDITIONAL_ TERRACOTTA_ SERVERS_ LIST=""	Comma-separated list of additional machine names hosting Clustering Service instances (for Clustering Service high availability).

Note: All the settings that you specify are the same for each Central node.

The following is an example of providing HA for the Clustering Service with one active node and two passive nodes. Note that each node that you add to the Clustering Service cluster slows system throughput, so that you need to keep system performance when deciding how many nodes to add. For more information on the impact of increasing the number of stand-by nodes, see ["Clustering Service Overview" \(on page 7\)](#).

The following is a sample of a Central server node installation


```
ICONCLUDE_HOME="Central-9.00"

# the location where OO Cluster is installed. Make sure that you points to the C
lustering's home folder
ICONCLUDE_CLUSTER_HOME="Cluster-9.00"

# Installation types:
# "terracotta_central" --> Configure OO node in Clustering Mode
# "terracotta_server" --> Install Clustering Service
# "load_balancer" --> Install Cluster Load Balancer

INSTALL_TYPE="terracotta_central"

#Terracotta server host name or IP address (do not use localhost, 127.0.0.1, etc
)
TS_SERVER="vancouver.mydomain.com"

# The path for the clustered Repository
SHARED_REPO_PATH="/opt/sharedRepo"

# List of comma-delimited Clustering Service machine addresses if exist
ADDITIONAL_TERRACOTTA_SERVERS_LIST="tacoma.mydomain.com,portland.mydomain.com"
```

5. Save your changes and close the `install.config` file.
6. Run the `configure.sh` script.
7. To restart the Central Service, run the following commands:

```
./bin/Central.sh stop
./bin/Central.sh start
```
8. **Important!** Verify that the Central and Clustering Service clusters are correctly configured. For details, see ["Verify That the Central Cluster is Correctly Configured" \(on page 36\)](#).
9. Connect the load balancer to the Central. For details, see ["Configuring a Load Balancing Cluster" \(on page 28\)](#).
10. Configure the Scheduler in the Central. For details, see ["Working with the Scheduler in a Cluster" \(on page 33\)](#).

Chapter 6

Load Balancer Configuration

This chapter includes:

Providing Load Balancing for OO Components	26
Providing High Availability to the Load Balancer.....	27
Installing the HP OO Load Balancer on a Windows Platform	27
Installing the HP OO Load Balancer on a Linux Platform	28
Configuring a Load Balancing Cluster	28
Connecting Studio to the Load Balancer	32
Upgrading the Load Balancer	32

Providing Load Balancing for OO Components

You can provide load-balancing for Central by installing the HP OO Load Balancer and then running the **cluster-config.jar** file (a Java command-line utility) to add Central and RAS, instances to the Load Balancer configuration.

Note: When using a load balancer, make sure that any interaction that had originally pointed to the Central now points to the load balancer. It can be a patch installation, flow execution, the DynamicallyLaunchFlow step, and so forth.

The OO Load Balancer is based on the Apache Load Balancer, version 2.2.11. The default configuration of the Apache load balancer that comes with OO is to spread the flows evenly between the various centrals based on the number of currently running requests. This is called **Pending Request Counting** (see [Apache documentation](#)). In addition, subsequent requests coming from the same web client are directed to the same Central using sticky sessions based on cookies.

This default OO behavior can be adjusted using the Apache load balancer configuration files. In addition, the customer can use a different load balancer with a different request spreading behavior.

For more information about the Apache Load Balancer, see the Apache Web site.

Load balancing is created by installing a load balancer, pointing it to the OO components, and integrating it with the Node Management in the Central. For information on creating the load-balancer-managing flow and enabling load-balancing management, see the HP Operations Orchestration Central User Guide.

1. Install the HP OO Load Balancer, using one of the two following procedures:
 - ["Installing the HP OO Load Balancer on a Windows Platform" \(on page 27\)](#)
 - ["Installing the HP OO Load Balancer on a Linux Platform" \(on page 28\)](#)

2. Configure the load balancer cluster. For details, see "[Configuring a Load Balancing Cluster](#)" (on [page 28](#)).
3. In Studio, create a flow that synchronizes the load balancer to the online/offline state changes of each node of OO Central.

For more information about building your own operations and flows, see the Guide to Authoring Operations Orchestration Flows and the Operations Orchestration Software Development Kit Guide.

Note: Operations and steps that point to the Central hostname should point to the load-balancing cluster's IP address.

4. In Central, enable load-balancer management for the Central cluster.

Note: When creating a load balancer management flow to manage the Apache load balancer's state (for details, see the section "Integrating a Load Balancer with Node Management" in HP Operations Orchestration Central User Guide), notice that the default value for the Apache load balancer management URL is **https://<lb-host>:<lb-port>/balancer-manager**.

If you use a software load balancer other than the OO Load Balancer, be sure that for the Central web client, the load balancer should be configured as a **sticky session**.

Note: Be sure to distinguish between creating a load-balancing cluster for RAS and Central instances. When you create a Central load-balancing cluster, add only Central servers. When you create a RAS load-balancing cluster, add only RAS instances.

Providing High Availability to the Load Balancer

To prevent the load balancer from becoming another "Single Point Of Failure", you must provide high availability for the load balancer. The load balancer can consist of two or more load balancer nodes that monitor each other using heartbeat. If one load balancer fails, the other takes over silently.

The common configuration for the load balancer is to have a virtual IP address that will be shared by the load balancer nodes and used for incoming HTTP requests.

Synchronizing between load balancer nodes can be configured using various 3rd party solutions. The relevant solution setup is out of this document's scope.

Installing the HP OO Load Balancer on a Windows Platform

Note: Before you perform the following steps, make sure jre 1.6 bin folder is in the PATH environment variable.

1. On the machine on which you are installing the Load Balancer, locate the **ClusterInstaller-9.00.exe** file on the HP OO Installation CD.
2. Double-click the file to open **Welcome** page of the Setup wizard.
3. On the **Welcome** page, click **Next**.

4. On the **License Agreement** page, accept the terms of the license agreement, and then click **Next**.

The **Installation Type** page appears.

5. To install the Load Balancer, select **Install Cluster Load Balancer**, and click **Next**.
6. On the **Select Destination Location** page, do one of the following, and click **Next**:
 - Accept the default location.
 - Type the location where you want the OO Load Balancer to be installed.
 - Click **Browse** and specify the required location.

The **Ready to Install** page appears, displaying the information that you specified.

7. To proceed, click **Install**. The Setup Wizard tracks progress on the **Installing** page.
8. When the installation is completed, click **Finish**.

Installing the HP OO Load Balancer on a Linux Platform

1. On the installation CD, in the Clustering subdirectory, locate and copy the **Cluster-9.00.zip** file and extract it to the location where you are going to perform the installation.

Note: The following steps create the information in the **INSTALL.txt** file (contained in the **Cluster-9.00.zip** file).

2. Open the **install.config** file for editing.
3. To specify installation of the Load Balancer as the install type, find the line that begins with `INSTALL_TYPE=` and make sure that the installation type is `"load_balancer"`.
4. Run the **configure.sh** script. The installation creates a service, **PASLB.sh**, in the `/apache` subdirectory of the Load Balancer home directory (which is, by default, `/Cluster-9.00`).
5. Using **configure.sh**, create or view the cluster and add or remove its nodes.

For information on adding or removing nodes, see "[Configuring a Load Balancing Cluster](#)" (on [page 28](#)).

6. To start the Load Balancer service, run the following command:

```
./apache/PASLB.sh start
```

Configuring a Load Balancing Cluster

Use the following procedure in both Windows and Linux environments.

Note:

- Be sure to distinguish between creating a load-balancing cluster for Central instances and RAS instances. When you create a Central load-balancing cluster (PAS), add only Central servers. When you create a load-balancing cluster for RAS installations, add only RAS instances.

- If you have multiple load balancers, make sure to perform the following procedure for each of them.
- The load balancer is already configured with a default entry called **Test Cluster** using port 8444. If you want to use this port, you must manually delete this entry before creating another one. If you are planning to use a port other than this one, you can ignore this entry.

To specify which OO components are load-balanced:

1. If you are using Windows:
 - a. Open a command window and Navigate to the location where you installed the HP OO Load Balancer and locate **cluster-config.jar**.

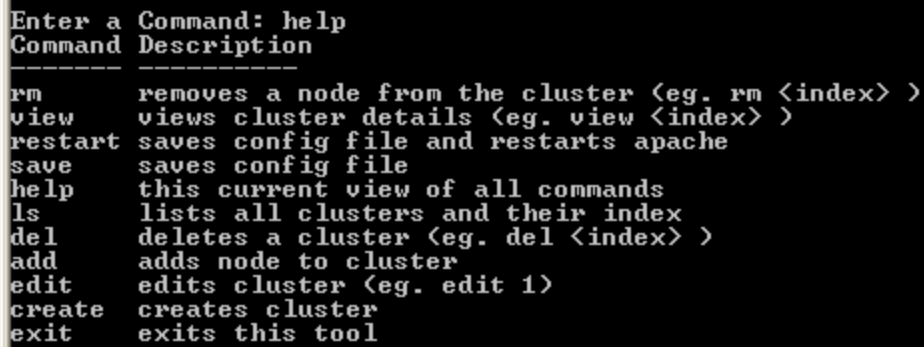
If you accepted the default locations when installing the Load Balancer, cluster-config.jar is in the OO home directory, in the \Clustering subdirectory.

- b. **Important! Before performing this step, make sure that Java 6 is installed.**

Run **cluster-config.jar**, using the following command:

```
java -jar cluster-config.jar
```

Note: To find the command help for **cluster-config.jar**, type **help**.



```
Enter a Command: help
Command Description
-----
rm      removes a node from the cluster (eg. rm <index> )
view    views cluster details (eg. view <index> )
restart saves config file and restarts apache
save    saves config file
help    this current view of all commands
ls      lists all clusters and their index
del     deletes a cluster (eg. del <index> )
add     adds node to cluster
edit    edits cluster (eg. edit 1)
create  creates cluster
exit    exits this tool
```

2. If you are using Linux:
 - a. Navigate to the location where you installed the HP OO Load Balancer and locate the extracted folder of the **Cluster-9.00.zip** file.
 - b. Run **configure.sh**.

Note: To find the command help for `configure.sh`, type `help`.

```

Enter a Command: help      this current view of all commands
Command Description
-----
help      this current view of all commands
restart   saves config file and restarts apache
save      saves config file
create    creates cluster
view      views cluster details (eg. view <index> )
edit      edits cluster (eg. edit 1)
del       deletes a cluster (eg. del <index> )
ls        lists all clusters and their index
rm        removes a node from the cluster (eg. rm <index> )
add       adds node to cluster
exit      exits this toolrestart saves config file and restarts apache

```

3. To create a cluster, type `create` and then press ENTER.
4. At the prompt **Enter the cluster name**, type the cluster name, and press ENTER.

Note: There should be two clusters – one for the Central instances and one for the RAS instances.

5. At the prompt **Enter the cluster type**, among the cluster types, do one of the following:
 - To add a Central node, type **PAS**
 - To add a RAS node, type **RAS**

Do not specify either JRAS or NRAS.
6. At the prompt **Enter the cluster port**, type an *unused* port for the cluster to use, and press ENTER.

The following message appears: **Cluster <clustername> successfully added!**

```

Enter a Command: create
Enter the cluster name: OO-PAS-Cluster
Enter the cluster type (PAS, RAS): PAS
Enter the cluster port: 8444
Cluster OO-PAS-Cluster successfully added!
Name: OO-PAS-Cluster
Host: PAS
Port: 8444
Nodes:
  Port      Host
  -----
(view) View Cluster, (edit) Edit Cluster, (add) Add Node, (rm #) Remove Node

```

7. To add a node to the cluster, type `add` and then press ENTER.
8. At the prompt **Enter the node host**, type the machine name of the node to add.
9. At the prompt **Enter the node port**, type the port that the node will use to communicate with the other nodes. The PAS default port is 8443 and the RAS default is 9004.

Each time you successfully add a node, the current state of the cluster, including its nodes, is displayed.

```

Enter a Command: add
Enter the node host: Central_1
Enter the node port: 8443
Node Central_1 successfully added!
Name: OO-PAS-Cluster
Host: PAS
Port: 8444
Nodes:
  Port      Host
  -----
  1         8443   Central_1

(view) View Cluster, (edit) Edit Cluster, (add) Add Node, (rm #) Remove Node

Enter a Command: add
Enter the node host: Central_2
Enter the node port: 8443
Node Central_2 successfully added!
Name: OO-PAS-Cluster
Host: PAS
Port: 8444
Nodes:
  Port      Host
  -----
  1         8443   Central_1
  2         8443   Central_2

(view) View Cluster, (edit) Edit Cluster, (add) Add Node, (rm #) Remove Node

```

10. To save the configuration file, at the prompt, type `save` and press ENTER.
11. To view the cluster details, type `view <index>`.

```

Enter a Command: view 1
Name: OO-PAS-Cluster
Host: PAS
Port: 8444
Nodes:
  Port      Host
  -----
  1         8443   Central_1
  2         8443   Central_2

```

12. To list all clusters and their index, press `ls`.

```

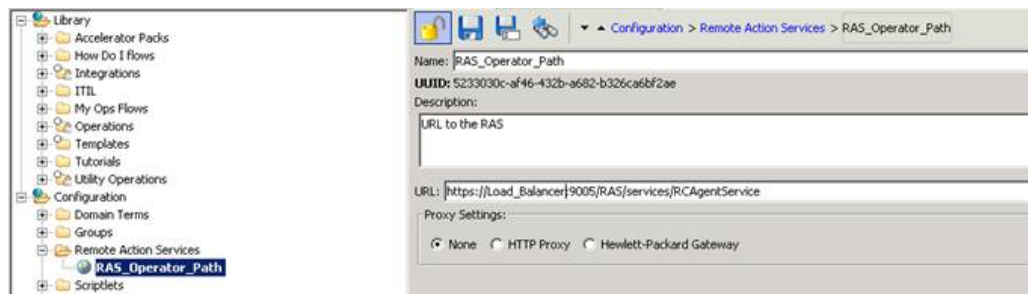
Enter a Command: ls
  Type      Port      Name
  -----
  1         PAS       8444    OO-PAS-Cluster
  2         RAS       9005    OO-RAS-Cluster

(view #) View Cluster, (create) Create Cluster, (del #) Delete Cluster

```

13. Type `exit`, press `y` and then press ENTER.
14. Restart the load balancer service (**RSCluster**).
15. To access RAS through the load balancer, RAS configuration needs to point to the RAS load balancer.
 - a. Open Studio.
 - b. Go to **Configuration > Remote Action Services**.
 - c. Edit the `RAS_Operator_Path` URL as follows: `https://<load balancer host name>:<load`

balancer RAS port number>/RAS/services/RCAgentService.



Connecting Studio to the Load Balancer

In a clustered environment, it is recommended to connect the Studio to the load balancer. This provides failover and load balancing capabilities to the Studio users when they access the Shared Repository.

Upgrading the Load Balancer

Upgrading the load balancer is achieved by uninstalling the existing load balancer and reinstalling the new load balancer version.

1. If you want to recreate the load-balancing cluster's membership after upgrading, record which computers are members of the load-balancing cluster.
2. Uninstall the Load Balancer.
3. Reinstall the Load Balancer, using the instructions in the appropriate one of the preceding sections:
 - ["Installing the HP OO Load Balancer on a Windows Platform" \(on page 27\)](#)
 - ["Installing the HP OO Load Balancer on a Linux Platform" \(on page 28\)](#)
4. If required, recreate the load-balancing membership by following the procedure in ["Configuring a Load Balancing Cluster" \(on page 28\)](#) for the components mentioned in the first step of this procedure.

Chapter 7

Scheduler Configuration

This section includes:

Introduction	33
Working with the Scheduler in a Cluster	33
Create a Clustered Scheduler.....	33
Configure the Scheduler to Work With a Load Balancer.....	34

Introduction

OO Scheduler is implemented using the Quartz scheduler. The Quartz scheduler provides its own cluster for high availability. Each Quartz scheduler instance can handle the execution of all the pre-scheduled jobs ensuring that if a certain scheduler instance is down, the others can take over and run the scheduled flows.

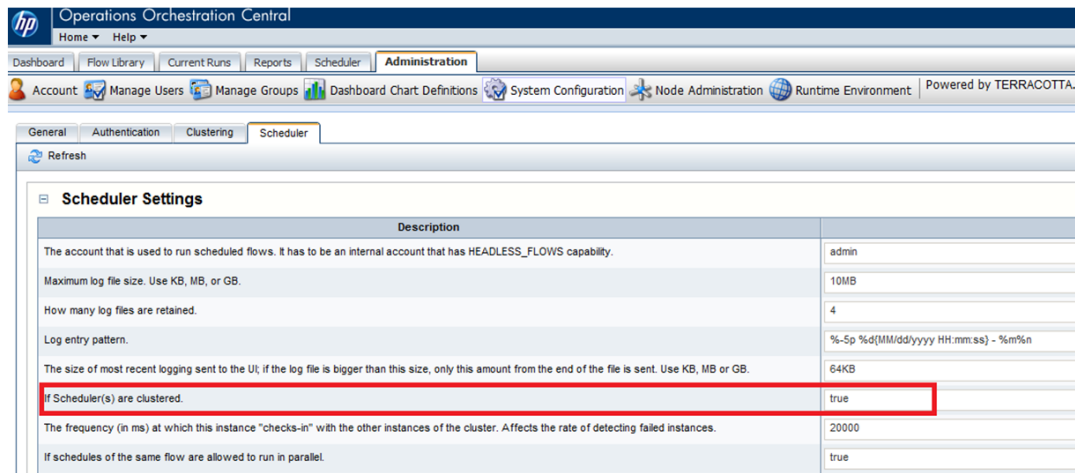
The various scheduler instances are constantly monitoring the pool of pre-scheduled jobs. Once a certain job is scheduled to run, the first scheduler instance marks the job as its own and then executes it. If the scheduler instance fails to execute the job, another scheduler instance takes over.

Working with the Scheduler in a Cluster

This section describes how to create a Quartz cluster and configure the scheduler to work with the load balancer.


Create a Clustered Scheduler

1. In HP Operations Orchestration Central, select **Administration > System Configuration > Scheduler**.
2. In the Scheduler Settings pane, type `true` in the **If Scheduler(s) are clustered** box.



When you have a Quartz cluster, it is enough to schedule the flow on only one of the cluster's nodes.

The flow you schedule on one node automatically appears on the other node(s). For example, Cluster1 contains two nodes —Node **A** and Node **B**. If you schedule a flow on Node **A**, as displayed in the picture below, the Quartz cluster automatically schedules the flow in the **Scheduler** tab on the second node as well.



If the scheduler service fails on Node **A**, the job is triggered on Node **B** by its scheduler, and is executed by its Central service. But, if the Central service fails on Node **A**, the flow is triggered by its scheduler but is not executed since the Central service is down. As a result, scheduled flows are not executed. To resolve this issue, it is recommended to configure the scheduler to work with a load balancer, as described in "[Configure the Scheduler to Work With a Load Balancer](#)" (on page 34).

Note: Each Central connects to its local scheduler service. If a certain scheduler service is down, then the Central on that node cannot schedule flow executions.


Configure the Scheduler to Work With a Load Balancer

Configuring the scheduler to work with a load balancer, ensures that all scheduled flows are executed in the Central cluster. The scheduler triggers the flow, and sends a request to the load balancer to execute the flow. The load balancer then sends the request to one of the Central Services. If the Central Service is down on one of the nodes, the load balancer submits the request to another available Central Service node.

To ensure that all scheduled flows are executed in a Central cluster, do the following:

1. On each scheduler, edit the parameters in the **<OO installation directory>\scheduler\conf\scheduler.properties** file as follows:
 - **dharmarepaircentral.host**=Load Balancer name or IP address
 - **dharmarepaircentral.port**=Load Balancer port
 - **dharmarepaircentral.proto**=Load Balancer protocol

For example:



2. Restart each scheduler service, one at a time.

Chapter 8

Common Tasks

This chapter includes:

Prepare the Shared File System	35
Verify That the Central Cluster is Correctly Configured	36

Prepare the Shared File System

All the nodes in a Central cluster share the repository. The repository is stored as a set of folders and files. This requires that the shared repository reside on a shared file system/shared folder, using the technology of your choice.

The Central cluster must be able to access this shared file system over the network.

- For a Windows-based Central cluster, you can create a regular Windows share and restrict access to it by using standard Windows file system security (using domain accounts).
- For a Linux-based Central cluster, you can use NFS, with proper access control.

Caution: If you use NFS for the shared repository, it is strongly recommended that you use NFS version 4 and turn off NFS caching.

The accounts under which Central will run must have full Read and Write permissions to the shared network directory (Windows or Linux).

- For a Windows-based Central cluster using a regular Windows share, switch the Central service to a domain account that has full permissions for the shared network directory.
- For a Linux-based Central cluster using NFS, assure regular Linux access control for the Central processes.

To achieve high availability for the HP OO cluster, the shared file system should be highly available as well. HP OO does not depend on a particular technology, because the repository is accessed through standard file IO.

The following procedure describes how to accomplish this on a Linux operating system.

To create and assign rights to a shared folder using the NFS Server on Linux:

1. Log on as the root user.
2. Check that NFS service is installed in each of the nodes and NFS server.
3. On each of the nodes, insert an NFS configuration such as the following in the file `/etc/fstab`.

```
<hostname>:/Linux_share /Linux_share nfs  
rsize=8192,wsize=8192,timeo=14,intr
```

where `<hostname>` is the host machine.

4. To give the mounted folder direct communication, insert the following line in the **/etc/exports** file on the NFS server: `/Linux_share 16.0.0.0/8(rw) .`
5. Mount the NFS folder in each of the nodes, using a command such as the following.

```
mount -t nfs <hostname>:/Linux_share /Linux_share
```

Verify That the Central Cluster is Correctly Configured

Note: Before starting Central, verify that the shared repository is accessible. Central does not start if the shared repository is down.

To verify that the high-availability Central cluster setup is correct, check the following on each Central node:

- In the Central home directory's `\conf` subdirectory, in the **Central.properties** file, find the line beginning with:

```
dharmarepository=
```

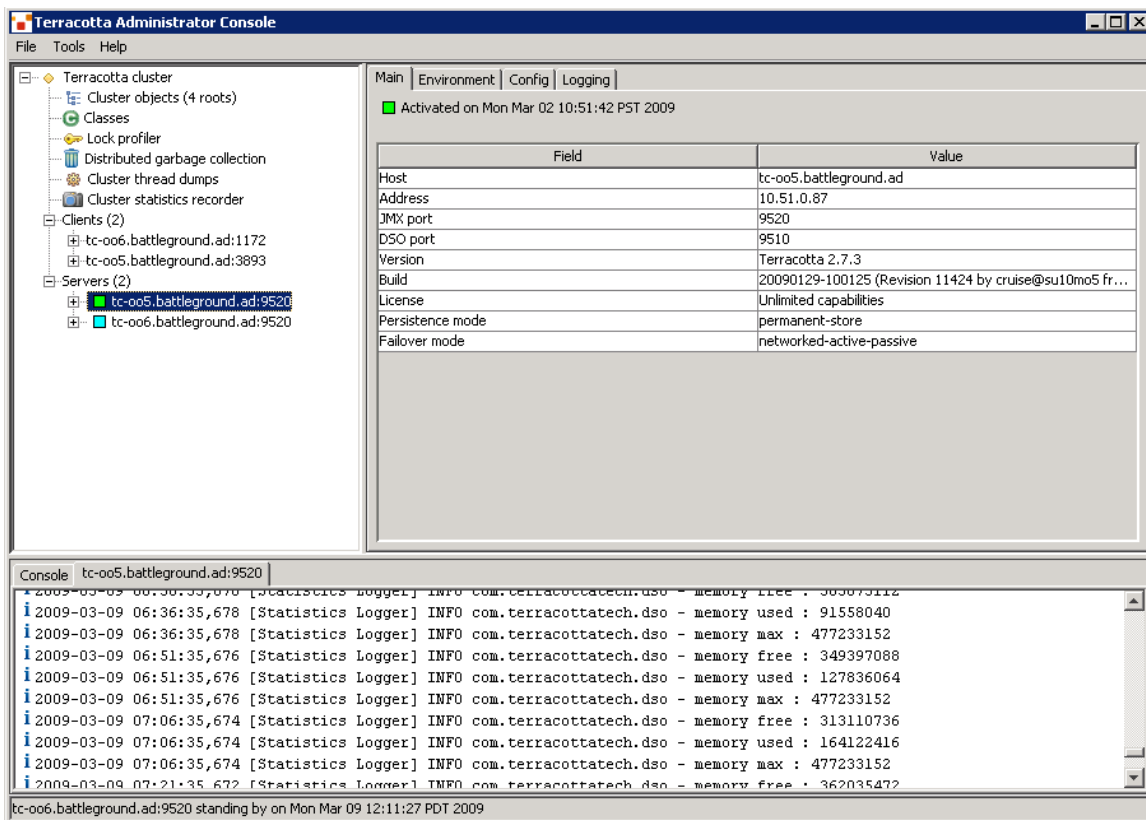
- The `location=` setting must point to the shared repository.

The reference to the shared repository can be either a UNC or a shared-directory reference. Whichever method of reference this setting uses, the reference must be consistent across the nodes of the Central. That is, the shared-repository reference must be identical in the `Central.properties` file on each node of the Central cluster. If the reference to the shared repository is not consistent across the Central cluster, the repository will be corrupted.

- The line should include the setting

```
shared=true;
```

- In the **<OO Cluster installation directory>\Cluster-9.00\tc-config.xml** file, confirm the Terracotta server settings and ports.
- In the **Clustering/terracotta/bin** folder, execute **admin.bat** on Windows or **admin.sh** on Linux, and connect to the cluster. An example of the minimal, highly available cluster is displayed below:



In the above example, two physical computers host both the Central Services (listed under **Clients (2)**) and the Clustering Service (listed under **Servers (2)**).

Under Clustering Services, make certain that one of the servers is green (green denotes an active server) and the other is blue (blue denotes a passive server).

- **If both servers are active, then do the following:**
 - In the <OO installation directory>\Clustering\terracotta\tc-config.xml file, verify that you used the machine names rather than the IP addresses.
 - If the file needs to be edited, restart the **RSGridserver services** after modifying the file one by one.
 - Run the <OO installation directory>\Operations **Orchestration\Clustering\terracotta\bin\Admin.bat** script and look for error messages in the Terracotta Administrator Console.
 - Look for error messages under: <OO installation directory>\Operations **Orchestration\Clustering\terracotta\logs\Terracotta_wrapper**.
- **If less than two Centrals appear under clients, then do the following:**
 - Check whether the repository is accessible. Central does not start if the shared repository is down.
 - Restart **Central Services**.

- Verify that there is one active and one passive server under the **Server** section in the Terracotta Administrator Console. If there is more than one active server, then follow the instructions in the section above explaining what to do if both servers are active.
- Check the logs in the following file: **<OO installation directory>\Operations Orchestration\Clustering\terraccotta\logs\central_warapper.log**.

Using this Terracotta Administrator Console is the most reliable way of visualizing the health of the cluster. Depending on your deployment topology, the number of services listed under both the **Servers** and **Clients** nodes may differ from the example above.

Chapter 9

Advanced Clustering Configuration

Removing a Node from an OO Cluster	39
Remove a Node from a Windows Cluster Installation.....	39
Remove a Node from a Linux Cluster Installation.....	40
Install a Standalone Clustering Service	41

Removing a Node from an OO Cluster

This section describes how to remove a node from an OO cluster and uninstall its components in Windows and Linux.

The steps include removing the node from the load balancer and removing the following node components — Clustering Service, Central, RAS, and the Scheduler.

Note: Once the node has been disconnected from the cluster, it is in invalid state and should not be used to execute or schedule flows.

This chapter includes:

Remove a Node from a Windows Cluster Installation	39
Remove a Node from a Linux Cluster Installation	40

Remove a Node from a Windows Cluster Installation

1. Remove the node from the load balancer.
 - a. Connect to the host on which the load balancer is installed.
 - b. From the **<Operations Orchestration Installation directory>\Clustering** folder, run the following script in a command prompt window: **java -jar cluster-config.jar**.
 - c. Type `ls` to view the list of clusters.
 - d. Type `view <cluster number>` to view the nodes of the selected cluster.
 - e. Type `rm <node number>` to remove the required node.
 - f. Type `exit`. When asked, **Save your changes?**, press **Y** and then ENTER.
2. Stop the clustering service.
 - a. Access the Windows Services window.
 - b. Locate the **RSGridServer** service.
 - c. Click **Stop Service**.
3. **Important! If you are using Windows Server 2003, make sure that the Services window is closed before performing this step.** Uninstall the clustering service.

- a. Access the Windows Programs and Features window.
- b. Locate **HP Operations Orchestration Cluster 9.02**.
- c. Click **Uninstall**.
4. Stop the HP Operations Orchestration Central Service.
 - a. Access the Windows Services window.
 - b. Locate the **RSCentral** service.
 - c. Click **Stop Service**.
5. Uninstall Central, RAS, and Scheduler. For details, see the section "Uninstalling OO Programs" in the Installing or Upgrading HP Operations Orchestration Guide.

Caution: After uninstalling the Clustering Service, the **RSGridServer** service, may not be removed. For details about how to uninstall the **RSGridServer** service, see "[Clustering Service Problems](#)" (on page 45).

Note: Uninstalling the clustering service invalidates HP-OO Central. To have a standalone Central installation, the Central service must be reinstalled. To reinstall Central, see the section "Installing or Upgrading Central" in the Installing or Upgrading HP Operations Orchestration Guide.

Remove a Node from a Linux Cluster Installation

Note: The commands in this section assume that OO is installed on the default path, that is, **/opt/hp**. If OO is installed elsewhere, substitute that path for **/opt/hp**.

1. Remove the node from the Load Balancer.
 - a. Connect to the host on which the load balancer is installed.
 - b. In shell, run the following command: **/opt/hp/<Operations Orchestration Installation directory>/Cluster-9.00/bin/configure.sh**.
 - c. Type **ls** to view the list of clusters and press ENTER.
 - d. Type **view <cluster number>** to view the nodes of the selected cluster and press ENTER.
 - e. Type **rm <node number>** to remove the required node and press ENTER.
 - f. Type **exit**. When asked, **Save your changes?**, press **Y**.and then ENTER.
2. Stop the Clustering Service.
 - a. Connect to the node of the OO Central, which you want to remove.
 - b. In shell,run the following command: **/opt/hp/<Operations Orchestration Installation directory>/Cluster-9.00/bin/cluster.sh stop**.
3. Stop the HP Operations Orchestration Central Service. In shell,run the following commands:
 - a. **/opt/hp/<Operations Orchestration Installation directory>/Cluster-9.00/bin/Central.sh stop**.

- b. `/opt/hp/<Operations Orchestration Installation directory>/Cluster-9.00/bin/Jras.sh stop.`
 - c. `/opt/hp/<Operations Orchestration Installation directory>/Cluster-9.00/bin/Scheduler.sh stop.`
4. In shell, run the following command: `rm /opt/hp/<HP Operations Orchestration Installation directory>.`

Install a Standalone Clustering Service

This procedure should be used when the Clustering Service instances are hosted on separate nodes from the Central instances.

Perform the following procedure on the machine that hosts the installed Clustering Service.

Caution: The installation path for the Clustering Service must not contain non-English language characters.

To install the Clustering Service separate from a Central server in Windows:

1. Locate the following executable file: **ClusterInstaller-9.00.exe**.
2. Double-click the file to open the Setup wizard.
3. On the **Welcome** page of the Setup wizard, click **Next**.
4. On the **License Agreement** page, accept the terms of the license agreement, and then click **Next**.

The **Installation Type** page appears, presenting the various types of installation and configuration options that you will use to create a Central cluster.

5. Select **Install Clustering Service**, and then click **Next**.
6. On the **Select Destination Location** page, type the location at which you want to install the Clustering Service files, and then click **Next**.

Or

Click the **Browse** button, navigate to the desired location, click **OK**, and then **Next**.

The **Clustering Service High Availability Information** page appears. This is the page on which you identify the other nodes that host Clustering Service instances.

7. Select the **Configure additional Clustering Services for High Availability** box.
8. In the text box, type a comma-delimited list of the machine names of the other Clustering Service nodes, and click **Next**.
9. On the **Ready to Install** page, click **Install**.
10. When the installation has completed, on the **Completing** page, click **Finish**.
11. **Recommended!** Review the troubleshooting issue regarding an active-active cluster.

To install the Clustering Service separate from a Central server in Linux:

Follow the steps in ["Add a Clustering Service to an Existing Central on Linux" \(on page 17\)](#).

Chapter 10

Troubleshooting

This section describes the troubleshooting issues for HP Operations Orchestration High Availability.

Split Brain Problem

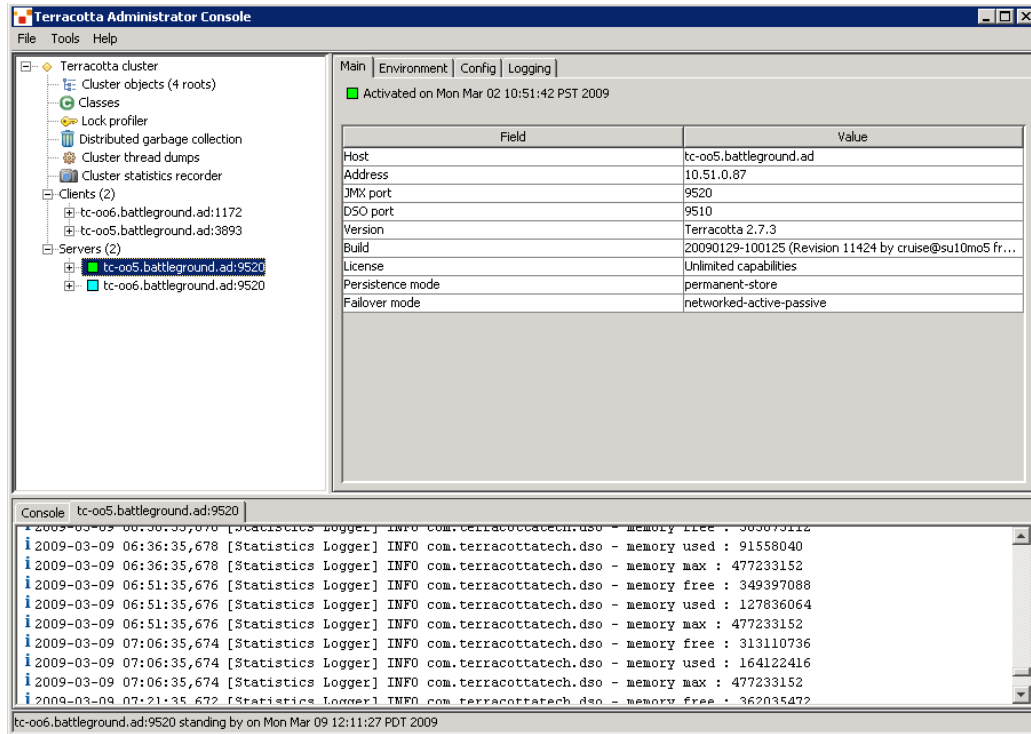
Problem: Check for the existence of a split-brain syndrome occurrence.

The split brain problem can occur when, due to configuration or network name resolution problems, the cluster contains more than one active Clustering Service (Terracotta server). If that happens, the correct functioning of the cluster is compromised, and the repository can become corrupted.

Note: If the cluster is not configured for HA (multiple Clustering Services), then the split brain syndrome is irrelevant.

Solution:

1. In Windows:
 - a. Open the Terracotta Administrator Console on one of the nodes that hosts the Clustering Service (Terracotta Server) as follows:
 - i. Change directories to **<installation_location>/Clustering/terracotta/bin**.
 - ii. Run **admin.bat**.
 - b. Connect to the localhost. A window similar to the Terracotta Administrator Console below should appear (if the Terracotta server on that node has correctly started). This window enables you to check for a split brain situation.



The list of Servers should contain only one active node. Click on each server to show its state. Only one should be shown as active (denoted by the green color) in the right pane (**Activated on <date and time>**).

2. In Linux:
 - a. Connect to each Clustering Service node.
 - b. Run **<Cluster Installation Directory>/Cluster-9.00/terracotta/bin/server-stat.sh**.

A result similar to the console below appears:

```
localhost.health: OK
localhost.role: ACTIVE
localhost.state: ACTIVE-COORDINATOR
localhost.jmxport: 9520
```

The **localhost.role** in only one of the nodes should have a value of **ACTIVE**.

3. If the cluster has a split brain, that is, two Clustering Services (under the **Servers** branch) appear as **Active**, reboot one of the Terracotta services and then verify if the problem has been resolved. If not, continue to the next step.
4. Reboot the entire cluster and examine the clustering-specific log files:

The clustering-specific log files are in the following locations:

- For Clustering Services (Terracotta servers):
 - **<installation_location>/Clustering/terracotta/logs**
 - **<installation_location>/Clustering/terracotta/terracotta-data/server-logs**

- For clustered Central services:
 - `<installation_location>/jetty/bin/terracotta-data/client-logs`
 - `<installation_location>/Central/logs`

Incorrect configuration of the list of Terracotta servers in the `<installation_location>/Clustering/terracotta/tc-config.xml` file is a common source of split brain syndrome.

In the `tc-config.xml` file on all nodes (either clustered Central, Clustering Service, or co-located Central and Clustering Services), the list of servers should be consistent, and the only difference should be the order in which Terracotta servers are listed. Every node that hosts the Clustering Service (Terracotta server) should have its first entry pointing to itself (do not use "localhost") and then the other additional Clustering Service machine names.

It is recommended that fully qualified names be used consistently across all nodes. If for some reason IP addresses have to be used on some nodes, it is recommended that all `tc-config.xml` on all cluster nodes be switched to use IP addresses.

Note: For each node that hosts the Clustering Service, the name attribute of the first server entry in `tc-config.xml` (should be pointing to either that machine's fully-qualified name or IP address) has to match the following entries in `<installation_location>/Clustering/terracotta/wrapper.conf` (or the Clustering Service will not start):

```
wrapper.app.parameter.4=<name attribute from tc-config.xml>
```

```
wrapper.app.parameter.9=<name attribute from tc-config.xml>
```

For example, if the first entry in `tc-config.xml`'s `<servers>` section for a node that hosts the Clustering Service has:

```
<server host="TC-006.battleground.ad" name="TC-006.battleground.ad">
```

...

then that node's `wrapper.conf` (at `<installation_location>/Clustering/terracotta/`) should contain:

```
wrapper.app.parameter.4=TC-006.battleground.ad
```

```
wrapper.app.parameter.9=TC-006.battleground.ad
```

If another node was configured (for instance, `TC-oo5.mydomain.ad`) for high availability, then the list should also contain `TC-oo5.mydomain.ad` (note that from Clustering Service's point of view, these machine names are case sensitive):

```
<!-- ***** -->
```

```
<!-- When configuring multiple terracotta servers for HA, add additional -->
```

```
<!-- servers bellow this comment and before the <ha> section -->
```

```
<!-- ***** -->
```

```
<server host="TC-oo5.mydomain.ad" name="TC-oo5.mydomain.ad">
```

...

As stated above, if IP addresses have to be used, then it is recommended that they be used consistently across all **tc-config.xml** and **wrapper.conf** files.

Note: The server names in the **tc-config.xml** file are case sensitive. Make sure to use the exact same name and case each time.

For information about the clustering technology, its configuration, and mode of operation, see the [Terracotta](#) Web site.

Clustering Service Problems

Problem: The Clustering Service does not start.

Solution: Check if the installation path of the Clustering Service contains non-English characters. The installation path for the Clustering Services must not contain non-English language characters. If the installation path contains non-English character, you need to reinstall the Clustering Service in a path containing English characters only.

Problem: The **RSGridServer** service is not removed after uninstalling the Clustering Service.

Solution: If the **RSGridServer** service is not removed after uninstalling the Clustering Service, do the following:

1. Access the Registry Editor (**Start > Run > type regedit**).
2. Locate then delete the **RSGridServer** folder at:

HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services

3. Restart the machine

Problem: When adding additional Central instances to the cluster (for details, see "[Installing Additional Central Instances in Windows](#)" (on page 22) you cannot use the IP address of the machines you want to use to host Clustering Services.

Solution: If you need to replace occurrences of IP addresses for the machines you want to use to host Clustering Services, perform the following steps:

1. Edit the **<HP Operations Orchestration installation directory>\Clustering\terracotta\tc-config.xml** file on all cluster machines.
2. Replace all the occurrences of IP addresses with their respective hostnames.
3. Save the file.
4. Access the Windows Services window and, one by one, restart the **RSGridServer** service on all cluster machines.

Problem: When attempting to restart OO processes on one node, you start the cluster first, and then Central, Central fails to restart. The following exception is seen in **Central_wrapper.log**:

```
Failed To Modify Field:  
com.iconclude.dharma.commons.repo.DharmaPathReference.path in  
com.iconclude.dharma.commons.repo.DharmaPathReference
```

Solution:

1. Edit `<OO installation directory>/Cluster-9.00/terracotta/tc-config.xml`.
2. Locate the **locks** section.
3. After the first **<autolock>** block, append these lines:

```
<autolock auto-synchronized="true">  
  
  <method-expression>*  
  com.iconclude.dharma.commons.repo.DharmaPathReference.*(..)</method-  
  expression>  
  
  <lock-level>write</lock-level>  
  
</autolock>
```

Terracotta Dirty Database Problem

Problem: The `<OO Installation directory>\Clustering\terracotta\logs\Terracotta_wrapper.log` contains the following error: `Error Type: Newly Joined Node Contains dirty database.`

This may cause the following problems:

- More than one active node is displayed in the Terracotta Administrator Console.
- Flow execution fails and the following error message appears in the Central UI: `RAS-side exception occurred`

Solution:

Perform the following steps on each cluster node:

1. Open the `<OO Installation directory>\Clustering\terracotta\wrapper.conf` and add the following line in the **Java Additional Parameters** section:

```
wrapper.java.additional.9=-Dcom.tc.l2.nha.dirtydb.autoDelete=true
```

2. Restart the RSGridServer service.

NFS Problem

Problem: When working with an NFS-based shared repository, the following message appears `java.io.IOException: Stale NFS file handle. when opening the Flow Library in the Web UI`. This can happen in the following cases:

1. When you start more than one Central using the same NFS repository simultaneously.
2. When the same user is logged in to the same Central in parallel sessions.

Solution:

Identify the Central on which the exception occurred and then restart it as follows:

1. Log in directly to the Web UI of each Central instead of using the load balancer URL.
2. Go to the **Flow Library** tab.
3. If the exception mentioned above appears, then restart this Central.

Corrupted Repository Problem

Problem: The Central Service fails due to a corrupted repository.

Solution:

1. Verify that the problem has been caused by a corrupted repository by locating an error in the **Central_wrapper.log** file that indicates a repository exception.
2. Replace the corrupted repository files with the an uncorrupted backup version from the **backups** folder:
 - a. In the **<main repository>** folder, delete all the folders except the **backups** folder.
 - b. From the **backups** folder, extract the latest, uncorrupted .zip file to the **<main repository>** folder.

