

Operations Orchestration Software

Software Version: 7.50

Clustering and Load Balancing for HP Operations Orchestration 7.50

Document Release Date: March 2009
Software Release Date: March 2009



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 Notices

© Copyright 2009 Hewlett-Packard Development Company, L.P.

Trademark Notices

All marks mentioned in this document are the property of their respective owners.

Finding or updating documentation on the Web

Documentation enhancements are a continual project at Hewlett-Packard Software. You can obtain or update the HP OO documentation set and tutorials at any time from the HP Software Product Manuals web site. You will need an HP Passport to log in to the web site.

To obtain HP OO documentation and tutorials

1. Go to the HP Software Product Manuals web site (<http://support.openview.hp.com/selfsolve/manuals>).
2. Log in with your HP Passport user name and password.
OR
If you do not have an HP Passport, click **New users – please register** to create an HP Passport, then return to this page and log in.
If you need help getting an HP Passport, see your HP OO contact.
3. In the **Product** list box, scroll down to and select **Operations Orchestration**.
4. In the **Product Version** list, click the version of the manuals that you're interested in.
5. In the **Operating System** list, click the relevant operating system.
6. Click the **Search** button.
7. In the **Results** list, click the link for the file that you want.

Where to Find Help, Tutorials, and More

The HP Operations Orchestration software (HP OO) documentation set is made up of the following:

- Help for Central
Central Help provides information to the following:
 - Finding and running flows
 - For HP OO administrators, configuring the functioning of HP OO
 - Generating and viewing the information available from the outcomes of flow runsThe Central Help system is also available as a PDF document in the HP OO home directory, in the \Central\docs subdirectory.
- Help for Studio
Studio Help instructs flow authors at varying levels of programming ability.
The Studio Help system is also available as a PDF document in the HP OO home directory, in the \Studio\docs subdirectory.
- Animated tutorials for Central and Studio
HP OO tutorials can each be completed in less than half an hour and provide basic instruction on the following:
 - In Central, finding, running, and viewing information from flows
 - In Studio, modifying flowsThe tutorials are available in the Central and Studio subdirectories of the HP OO home directory.
- Self-documentation for operations and flows in the Accelerator Packs and ITIL folders
Self-documentation is available in the descriptions of the operations and steps that are included in the flows.

Support

For support information, including patches, troubleshooting aids, support contract management, product manuals and more, visit the following site:

- <http://support.openview.hp.com>

Contents

Warranty	ii
Restricted Rights Legend	ii
Trademark Notices	ii
Finding or updating documentation on the Web.....	iii
Where to Find Help, Tutorials, and More	iii
Support	iv
Clusters and HP OO: an overview	1
Failover and horizontal scalability of Central servers	1
Clustering Service overview	1
Load-balancing Central servers and any standalone RAS installations.....	2
Simple deployment configuration.....	2
Scaling out the cluster by adding more Central nodes	4
Scaling out the cluster by adding Central nodes and separating them from Clustering Service instances.....	5
If you are upgrading	6
Overall procedure for clustering and load-balancing Central	6
Tasks: Creating an HP OO cluster that co-locates Central and Clustering Service instances.....	7
Tasks: Creating an HP OO cluster that separates Central and Clustering Service instances.....	8
Preparing the shared file system	8
Creating an HP OO cluster on Windows systems.....	9
Co-locating the Clustering Service and Central service on Windows	9
Installing a Clustering Service on Windows	11

Configure a Central Service instance to run in clustering mode on Windows.....	13
Creating an HP OO Cluster on Linux systems.....	15
Installing a Clustering Service on Linux.....	15
Configuring a Central Service instance to run in clustering mode on Linux	18
Verifying that the Central cluster is correctly configured.....	20
Load-balancing for Central and RASes	21
Installing Load Balancer on Windows	22
Installing Load Balancer on Linux.....	23
Creating a load-balancing cluster and adding components to load-balance.....	23
Connecting Studio to the Central shared repository through the Load Balancer.....	24

Troubleshooting the HP OO Cluster deployment	25
--	----

Index	28
-------------	----

Clusters and HP OO: an overview

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, and automatic recovery of failed flows when a Central node fails. HP OO Load Balancer provides load balancing for requests to the Central servers and standalone RAS installations.

Failover and horizontal scalability of Central servers

- Central servers use an active-active topology, thus maximizing the utilization of available hardware.
- 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 an 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).

If a node in the Central cluster fails:

- The failed Central node's flow runs are automatically recovered on another node of the cluster and resumed at their last check-pointed step. This feature is not enabled by default, but it can be enabled by an OO administrator (please see the *OO Administrator's Guide*). Also note that only unassisted (headless) flow runs can be recovered by this feature.
- 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 user-private workspace. Note that the system maintains edits based on explicit saves of the ongoing work; when the connection to a Central node is lost, the Studio un-saved edits cannot be recovered.

Clustering Service overview

HP Operation Orchestration's clustering is centered on the concept of a clustering service that is the cluster's arbiter and controls all aspect related to clustered state and coordination. The Clustering Service is a separated 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 so that there is no contention for resources between a Central server and a Clustering Service server.

The clustering support is based on a Java VM-level clustering solution called Terracotta (www.terracotta.org). The OO Clustering Service is a Terracotta server in either active or hot-standby mode. Note that Central processes themselves use active/active topology. Only the Clustering Services use an active/hot-standby topology.

The deployment of a hot-standby Clustering Service is not strictly required. All that is required for the proper functioning of an HP OO cluster is the existence of the active Clustering Service. The standby service is recommended for achieving high availability in case of the active service's failure. 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 standby Clustering Services will decrease throughput. For most deployments, one active and one standby Clustering Service are sufficient.



Best Practice: It is recommended that both the active and the standby Clustering Services run on identical hardware.

Load-balancing Central servers and any standalone RAS installations

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.



Best Practice: It is recommended that you provide a highly available load balancer solution, as this component can become a single point of failure.

The Central database contains the Central nodes' run histories, schedules, and system configuration.



Best Practice: It is recommended that you provide a highly available database solution, as this component can become a single point of failure.

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

Simple deployment configuration

The simplest deployment configuration that is highly available (from Central's point of view) is a two-node Central cluster, in which each node hosts one Central instance and one Clustering Service instance, on identical hardware. When the cluster starts, the Clustering Services go through an election process and one of them becomes the active instance. Full high availability requires the following:

- Highly available database
- Highly available shared file system (for hosting the shared repository)

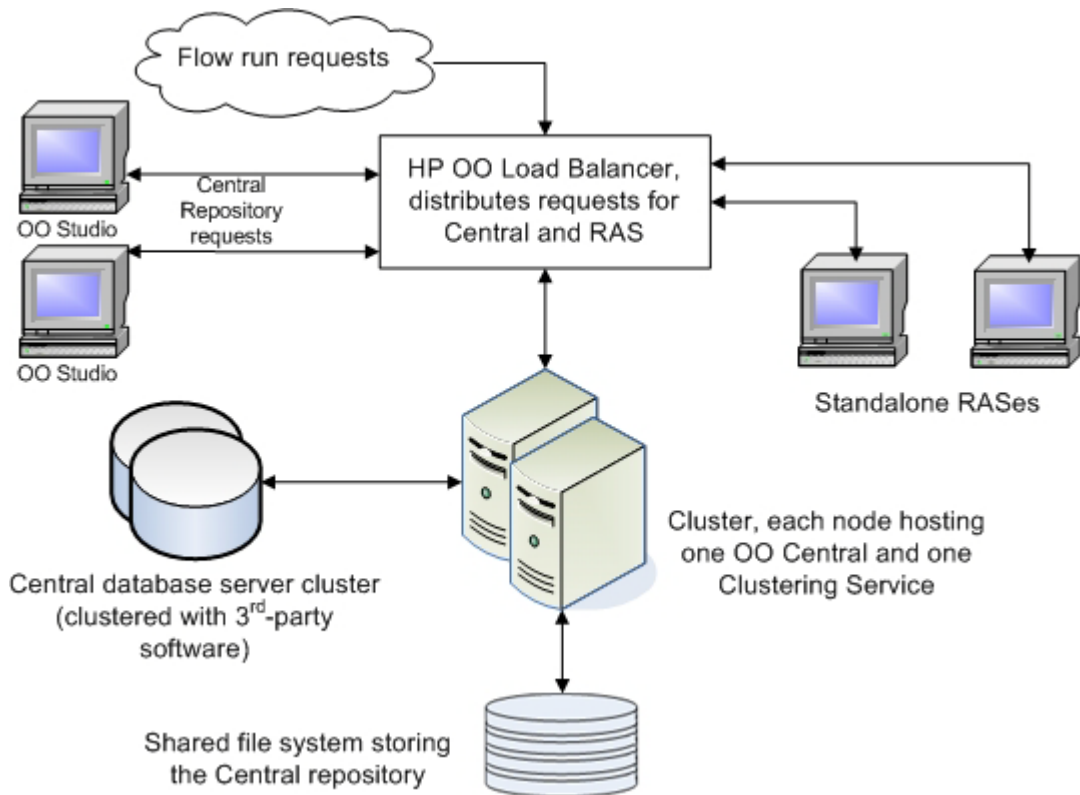


Figure 1 - Load balancing and high-availability clustering when the Clustering Service and Central Service are co-located

- If one node fails: if the failed node hosted the active Clustering Service, the hot-standby becomes the active Clustering Service and the remaining node resumes functionality. If the failed node comes back up, the Clustering Service will become hot-standby, and its Central instance will join the cluster and be ready for performing work.
- If the Clustering Service fails on one node: if the failed Clustering Service was the active instance, the hot-standby takes over and becomes active. The Central instances resume functionality after this failover is completed. If the failed Clustering Service was the standby, the cluster continues functioning, without a standby. Further failure of the active instance results in cluster "death" and requires the restart of all nodes.
- If the Central service fails on one node: the other Central continues functioning, and when the failed Central instance comes back up, it joins the cluster and is ready for work. As mentioned above, Studio and Central UI sessions are not failed over when one Central service instance fails. Headless flow runs can be automatically resumed on the remaining instance if the feature is enabled.

Scaling out the cluster by adding more Central nodes

The simple deployment presented above 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.

Important: Multiple clustered Central processes running on the same node is not supported. While 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.

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

For example, a scaled out cluster of five nodes could have three nodes that host both Central and Clustering Service instances (so there are two stand-by Clustering Services allowing failures of up to two nodes hosting a Clustering Service before the HP OO cluster is non-functional) and two nodes that only host Central instances.

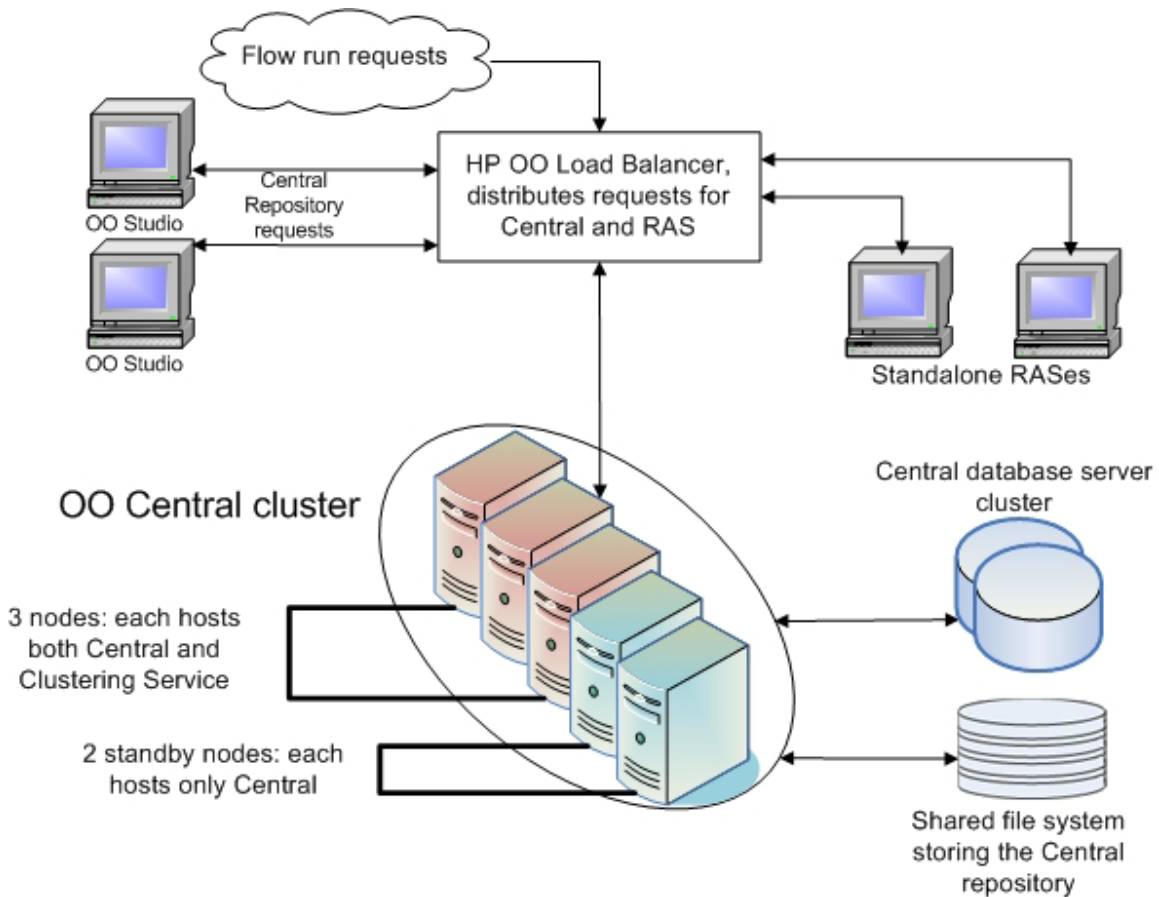


Figure 2 - Five-node Central cluster with two standby nodes for Clustering Service

Scaling out the cluster by adding Central nodes and separating them from Clustering Service instances

The above deployment scenarios have a common theme—the Clustering Services were co-located with Central instances. This has the side effect that the Clustering Service and the Central service compete for host resources. When the workload is high, this contention for resources can slow down the Clustering Services (the active one is especially susceptible to this) which results in decreased throughput. So, based on specific performance, scale requirements, and type of cluster (a development cluster, used mainly for developing flows, or a production cluster, mainly for running flows, or both), there may be cases when it is desirable for the Clustering Services to be hosted on separate nodes from the Central instances.

This deployment configuration results in an HP OO cluster containing two classes of nodes—Central nodes and Clustering Service nodes. As explained above, a large HP OO cluster can be serviced by a small number of Clustering Service nodes (a minimum of two are required for HA).

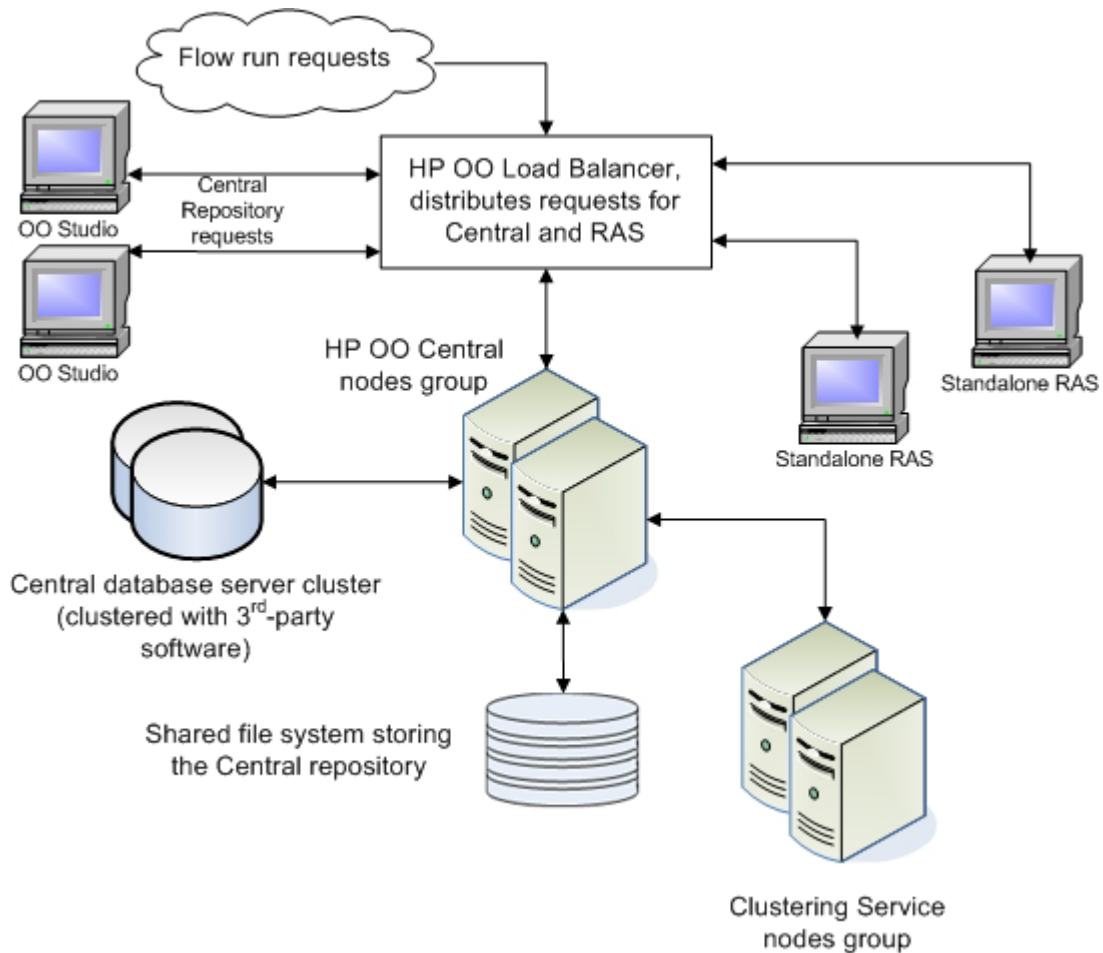


Figure 3 - Load balancing and high-availability clustering with distinct node groups for Central and Clustering Service

If you are upgrading

To upgrade the OO Load Balancer, uninstall the Load Balancer, then reinstall it, using the instructions in the following sections.

To upgrade a previous clustered OO deployment, you:

- Upgrade every individual Central node.
- Follow the procedures described below, where the Central installation step is replaced by the Central upgrade step.

Overall procedure for clustering and load-balancing Central

Creating a Central cluster involves:

- Creating a shared file system.



Best Practice: It is recommended that you not use any of the cluster nodes for hosting the shared file system.

- Installing Centrals on the desired nodes, all pointing to the same database, making sure that every Central instance is functional.
- Giving Read/Write file-system level access to Central processes to the share. This step depends on the chosen type of the shared file system.
- Creating the Central cluster by configuring Central on each Central node to run in cluster mode.
- Installing the OO Load Balancer.

How you carry out the above depends on which arrangement of Clustering Service and Central clusters you are setting up and on the server operating system on which you're installing.

- If you are co-locating Central and the Clustering Service instances, carry out the tasks in [Tasks: Creating an HP OO cluster that co-locates Central and Clustering Service instances](#).
- If you're creating the Central cluster on one set of servers and the Clustering Service cluster on another, carry out the tasks in [Tasks: Creating an HP OO cluster that separates Central and Clustering Service instances](#).

Tasks: Creating an HP OO cluster that co-locates Central and Clustering Service instances

1. Create the database cluster.

How you create the database cluster depends on the third-party software that you use create and support it.

2. Prepare the shared file system and give proper access to it (Central processes need Read/Write permissions to the share) see [Preparing the shared file system](#).

3. On each server:

- a. Install OO Central.

For information on installing Central, see [Installing or Upgrading HP Operations Orchestration \(InstallGuide.pdf\)](#).

- b. If the servers run on Windows, install and configure the Clustering Service and switch the Central instance to clustering mode, using [Co-locating the Clustering Service and Central service on Windows](#).

- c. If the servers run on Linux, install and configure the Clustering Service and switch the Central instance to clustering mode, using [Configuring a Central Service instance to run in clustering mode on Linux](#).

4. Install the HP OO Load Balancer (optional).

For installation and configuration procedures, see one of the following:

- [Installing Load Balancer on Windows](#) and [Creating a load-balancing cluster and adding components to load-balance](#)
- [Installing the Load Balancer on Linux](#) and [Creating a load-balancing cluster and adding components to load-balance](#)

Tasks: Creating an HP OO cluster that separates Central and Clustering Service instances

1. Create the database cluster.
How you create the database cluster depends on the third-party software that you use for creating and supporting the cluster.
2. Prepare the shared file system and give proper access to it (Central processes need Read/Write permissions to the share). See [Preparing the shared file system](#) for more information.
3. On each Clustering Service node, install and configure the Clustering Service.
For installation and configuration procedures, see one of the following:
 - [Installing a Clustering Service on Windows](#)
 - [Installing a Clustering Service on Linux](#)
4. Configure each Central node to run in clustering mode, using one of the following procedures:
 - [Configure a Central Service instance to run in clustering mode on Windows](#)
 - [Configuring a Central Service instance to run in clustering mode on Linux](#)
5. Install Load Balancer.
For installation and configuration procedures, see one of the following:
 - [Installing Load Balancer on Windows](#) and [Creating a load-balancing cluster and adding components to load-balance](#)
 - [Installing the Load Balancer on Linux](#) and [Creating a load-balancing cluster and adding components to load-balance](#)

Preparing the shared file system

All nodes in a Central cluster have to 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 share over the network. For a Windows-based Central cluster, a regular Windows share can be used, and the access to it can be restricted by using standard Windows file system security (using domain accounts). For a Linux-based Central cluster, NFS can be used, and proper access control should be employed.

Give the accounts under which Central will run full Read and Write permissions to the share. For a Windows-based Central cluster using a regular Windows share, you should switch the Central service to a domain account that has Read and Write access to the share. For a Linux-based Central cluster using NFS, assure regular Linux access control for the Central processes.

In order 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, as the repository is accessed through standard file IO. If the high availability requirements are less stringent, one could very well host the shared repository on the same infrastructure as the database.

Creating an HP OO cluster on Windows systems

Co-locating the Clustering Service and Central service on Windows

This procedure should be used when Central and Clustering Services are co-located on the same machine. This assumes that Central is already installed on this node. Run the following procedure on each such node.

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

1. On the HP OO Installation CD, navigate to and double-click the file **ClusterInstaller-7.50.exe**.
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 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.

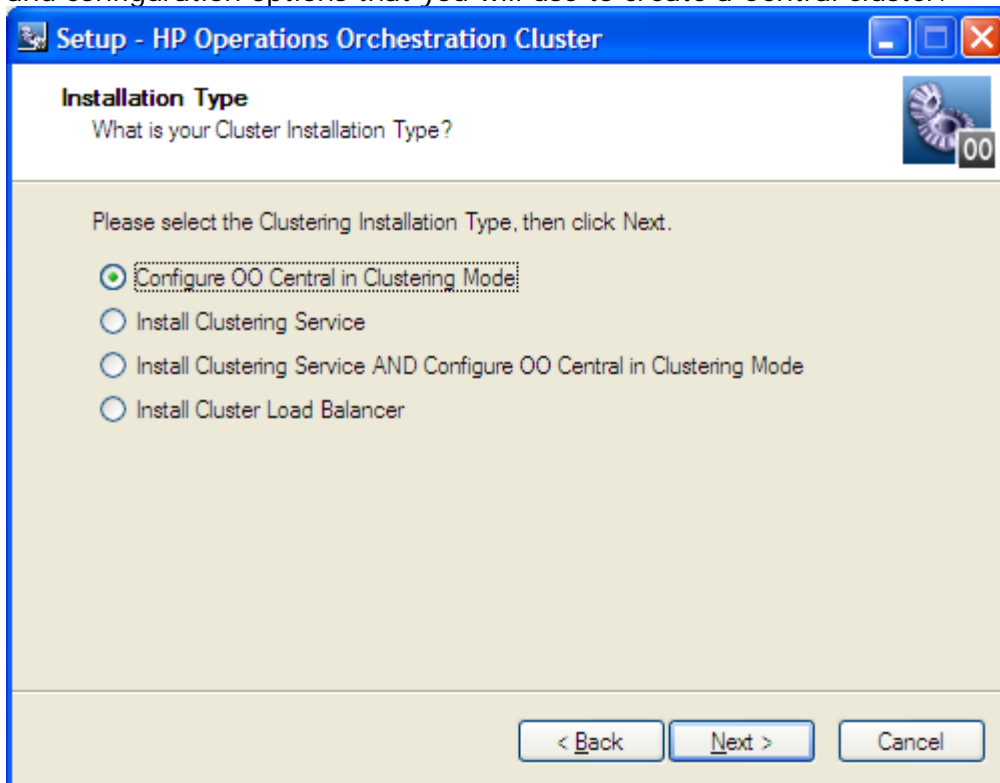


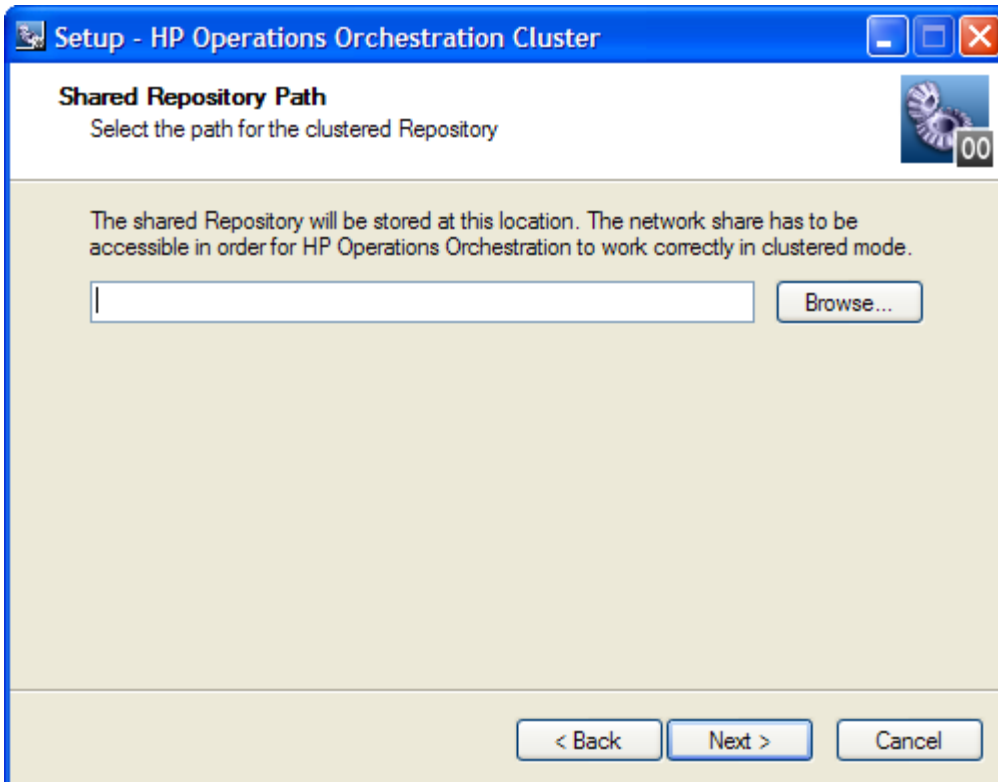
Figure 4 - Choosing the type of installation

4. Select **Install Clustering Service AND Configure OO Central in Clustering Mode**, and then click **Next**.
5. In the **Shared Repository Path** page's text box, type the location of the shared folder, and then click **Next**.

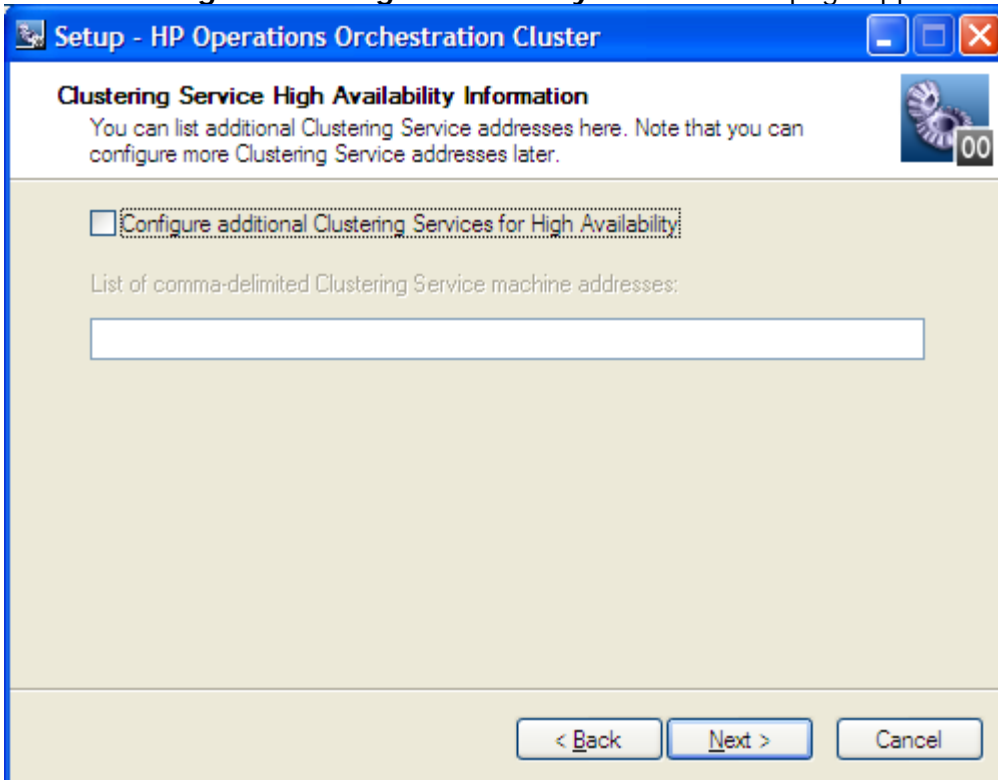
OR

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

The **Share Repository Path** page appears.



6. Enter or select the path for the shared repository, and then click **Next**.
The **Clustering Service High Availability Information** page appears.



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 in the cluster (that is, the other machines that will be hosting Clustering Service services).
9. Click Next.
10. On the **Ready to Install** page, click **Install**.
11. When the installation has completed, on the **Completing** page, click **Finish**.
12. **Important!** Verify that the Central cluster and Clustering Service cluster are correctly configured.

For information on verifying the configuration of the clusters, see [Verifying that the Central cluster is correctly configured](#).

To add the load-balancing capability to the Central cluster, you install the OO Load Balancer and list Central cluster nodes for which the Load Balancer should act. For information on installing and configuring the Load Balancer, see [Installing Load Balancer on Windows](#) and [Creating a load-balancing cluster and adding components to load-balance](#).

Installing a Clustering Service on Windows

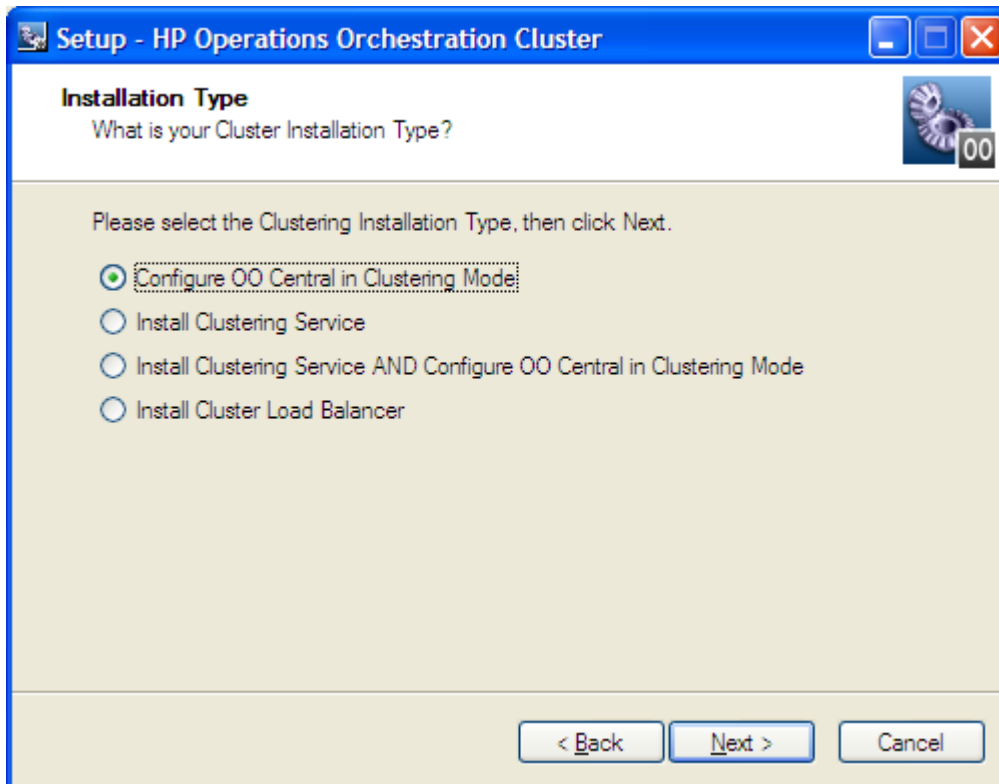
This procedure should be used when the Clustering Service instances are hosted on separate nodes from the Central instances.

Carry out the following procedure on each Clustering Service server.

To install the Clustering Service on a Windows system

1. On the HP OO Installation CD, navigate to and double-click the file **ClusterInstaller-7.50.exe**.
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 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. You will run this executable (ClusterInstaller-7.50.exe) several times in the course of creating a Central cluster and load balancing, choosing a different option on this page each time, depending on which task you're carrying out.

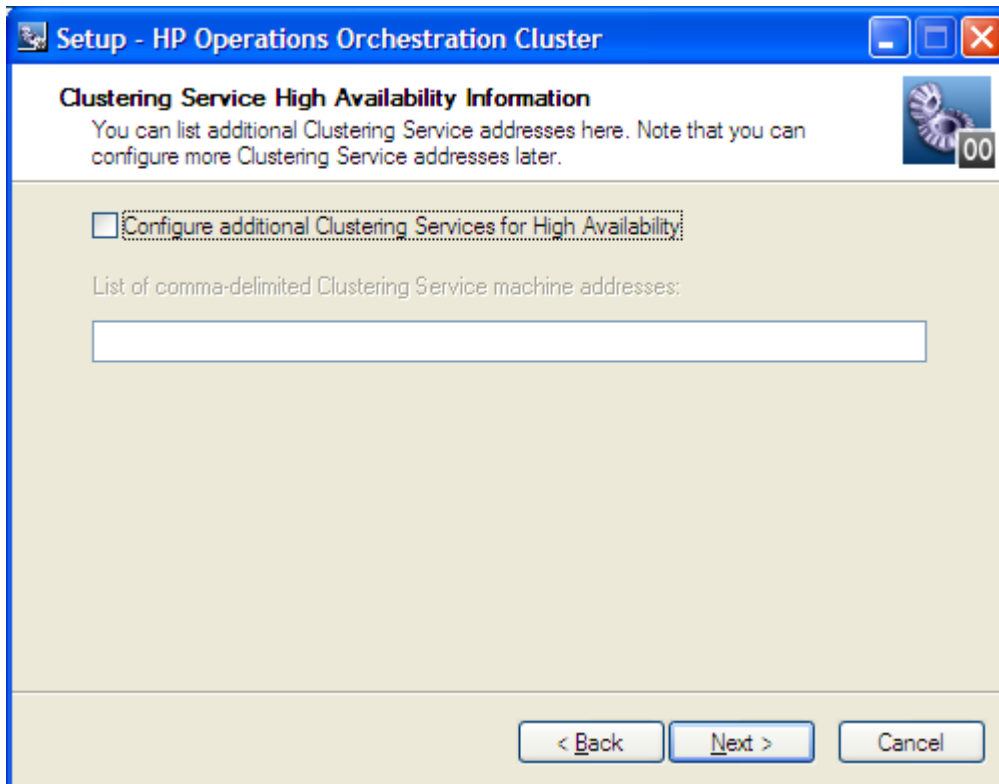


4. Select **Install Clustering Service**, and then click **Next**.
5. On the **Select Destination Location** page, type the location at which you want to install the Clustering Service files, and then click **Next**.

OR

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



6. Select the **Configure additional Clustering Services for High Availability** box.
7. In the text box, type a comma-delimited list of the machine names of the other Clustering Service nodes, and then click **Next**.
8. On the **Ready to Install** page, click **Install**.
9. When the installation has completed, on the **Completing** page, click **Finish**.

Configure a Central Service instance to run in clustering mode on Windows


This procedure should be used when additional Central instances are added to the cluster.

You create a Central cluster by configuring its nodes to run in clustering mode. Thus, to create a Central cluster, you carry out the procedure in this section on each Central cluster node.

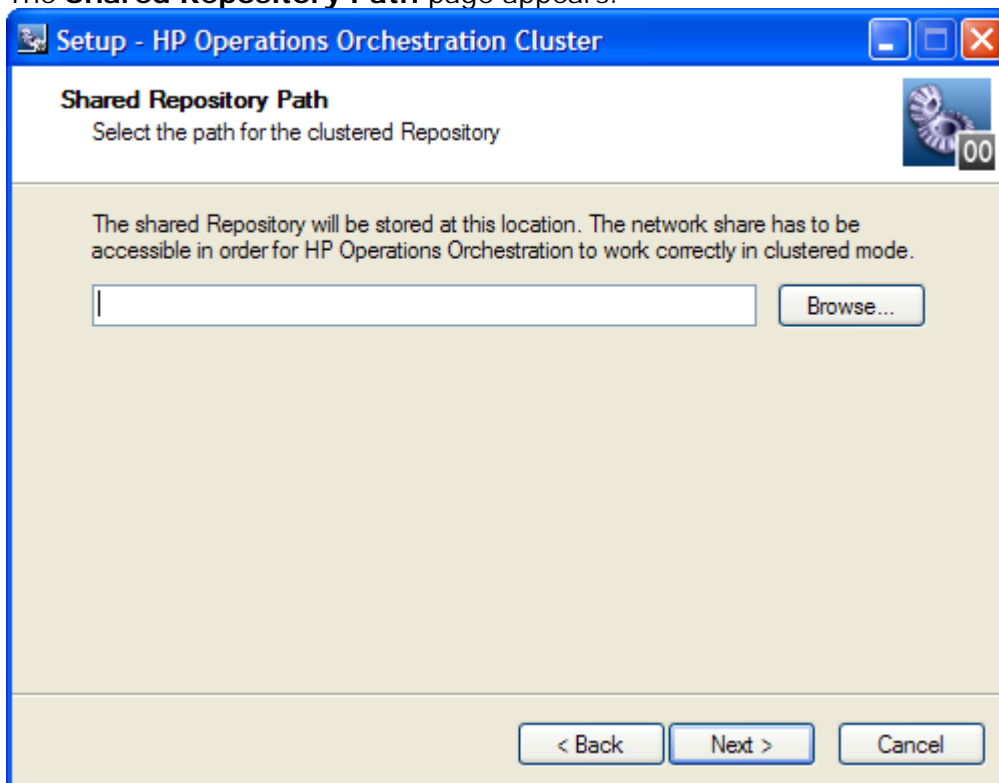
To create a Central cluster on Windows or add a node to an existing Central cluster

1. On the HP OO Installation CD, navigate to and double-click the file **ClusterInstaller-7.50.exe**.
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 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.



4. Select **Configure OO Central in Clustering Mode**, and then click **Next**. The **Shared Repository Path** page appears.

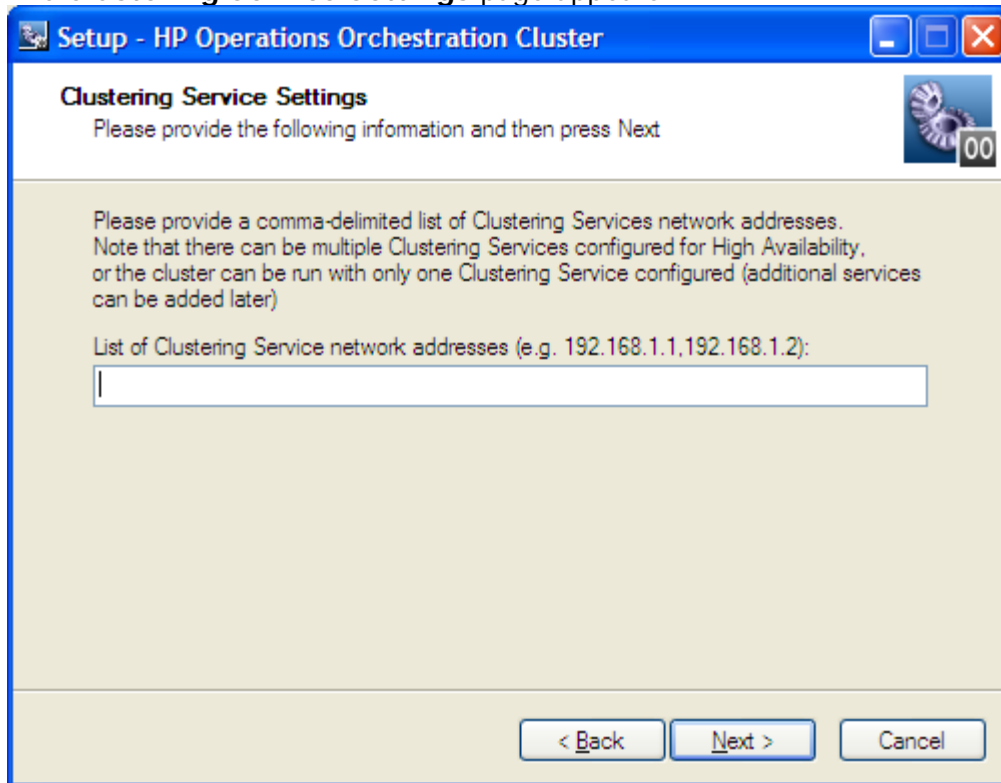


5. In the text box, type the location of the shared folder, and then click **Next**.

OR

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

The **Clustering Service Settings** page appears.



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

You can enter more than one machine name, if the cluster has multiple clustering Services (for supporting high availability of the Clustering Service itself).

7. On the **Ready to Install** page, click **Install**.
8. When the installation has completed, on the **Completing** page, click **Finish**.
9. **Important!** Verify that the Central cluster and Clustering Service cluster are correctly configured.

For information on verifying the configuration of the clusters, see [Verifying that the Central cluster is correctly configured](#).

Creating an HP OO Cluster on Linux systems

Installing a Clustering Service on Linux

The Clustering Service stores information that defines, maintains, and supports Central. If you have the available resources, it is strongly recommended that you configure additional Clustering Services.

Carry out the following procedure on each Clustering Service server or node.

To install the Clustering Service on a Linux system

1. On the installation CD, in the Clustering subdirectory, locate and copy the Cluster-7.50.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-7.50 directory.

Note: The following instructions reproduce the information in INSTALL.txt (contained in Cluster-7.50.zip).

2. In the Cluster-7.50 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 ("")
Clustering Service home directory	<code>ICONCLUDE_CLUSTER_HOME=""</code>	The location of the Clustering Service home directory (the location OS where Cluster-7.50.zip was extracted).
Installation of the Clustering Service and configuration of the server as a Clustering Service node	<code>INSTALL_TYPE=""</code>	<code>terracotta_server</code>
Clustering Service server	<code>TS_SERVER=""</code>	The machine name or IP address (on which you are installing the Clustering Service). The fully qualified machine name should be used, and should not be localhost.
Any additional Clustering Service servers	<code>ADDITIONAL_TERRACOTTA_SERVERS_LIST=""</code>	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 properly on vancouver.mydomain.com, it will look like the following.

```
# points to Central's home folder
ICONCLUDE_HOME=""

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

# 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=""

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

Figure 5 - Configuring installation of Clustering Service

3. Run the configure.sh script.
The configure.sh script creates CLUSTER.sh in the /bin folder.
4. To start the Clustering Service, run the following command:
`./bin/CLUSTER.sh start`
5. Repeat this procedure on each of the other nodes hosting Clustering Services.



Key information: 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 you're working on.

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"

Configuring a Central Service instance to run in clustering mode on Linux

To create the cluster, you make each Central server a Central cluster node by configuring it to run in cluster mode, with reference to the Clustering Service server (or servers) and the location of the shared repository. You do so by performing the following procedure on each Central server that is to be a node in the Central cluster.

To add Central servers that run Linux systems to the Central cluster, you use the following procedure, as you do when you first create the Central cluster. For information on installing Central on a Linux system, see *Installing or Upgrading HP Operations Orchestration* (InstallGuide.pdf).

The procedure assumes that you have installed Central on each Central server.

To add the server to the Central cluster

1. On the installation CD, in the Clustering subdirectory, locate and copy the Cluster-7.50.zip file and extract it to the machine which hosts the Central Service that is being configured in clustering mode.

Note: If the Central Service and Clustering Service are being co-located, Cluster-7.50.zip should only be extracted once on that node. The configuration script will be run twice, once for configuring the Clustering Service and once for configuring Central Service in clustering mode.

Extracting the .zip file creates the Cluster-7.50 directory.

Note: The following instructions reproduce the information in INSTALL.txt (contained in Cluster-7.50.zip).

2. In the Cluster-7.50 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	<code>ICONCLUDE_HOME=""</code>	The location of the Central home directory.
Clustering Service home directory	<code>ICONCLUDE_CLUSTER_HOME=""</code>	The location of the Clustering Service home directory (the location OS where Cluster-7.50.zip was extracted).
Configuration of the Central server as a Central node	<code>INSTALL_TYPE=""</code>	<code>terracotta_central</code>
Clustering Service server	<code>TS_SERVER=""</code>	Machine name of Clustering Service host server (first Clustering Service node if more are configured for HS).

To specify this	In this line of install.config	Type this between the quotation marks ("")
Repository shared by Central nodes	<code>SHARED_REPO_PATH= ""</code>	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	<code>ADDITIONAL_TERRACOTTA_SERVERS_LIST= ""</code>	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.

For example, you might specify the following in the install.config file:

- The Central home directory
/home/admin/Central-7.50
- The Clustering Service home directory
/home/admin/Cluster-7.50
- In `INSTALL_TYPE`, configuration of a Central node
`terracotta_central`
- The machine name or IP address of the Clustering Service server or the Clustering Service cluster
`TS_SERVER="vancouver.mydomain.com"`
- The location of the repository that the Central nodes share
/home/admin/mnt/pubflows
Note that the location is provided through a mount point (in this case, /mnt/) for the shared location.
- Comma-separated list of Clustering Service servers other than the one specified by `TS_SERVER= ""`

With the above values, the install.config file would look like the following:

This is an example of providing HA for the Clustering Service with one active node and two stand-by nodes. Note that adding standby nodes for the Clustering Service slows system throughput, so that there is a tradeoff between high availability and system performance. For more information on the impact of increasing the number of stand-by nodes, see [Clustering Service overview](#).

```

# points to Central's home folder
ICONCLUDE_HOME=""

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

# 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"

#Terracota 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="/home/admin/mnt/pubflows/"

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

```

Figure 6 - Sample configuration of a Central server node installation

3. Close the install.config file, saving your changes.
4. Run the configure.sh script.
5. To restart the Central Service, run the following commands:

```

./bin/Central.sh stop
./bin/Central.sh start

```

6. Repeat this procedure on each of the other Central nodes in the cluster.
7. **Important!** Verify that the Central cluster and Clustering Service cluster are correctly configured.

For information on verifying the configuration of the clusters, see [Verifying that the Central cluster is correctly configured](#).

To add the load-balancing capability to the Central cluster, you install the OO Load Balancer and list Central cluster nodes for which the Load Balancer should act. For information on installing and configuring the Load Balancer, see [Installing the Load Balancer on Linux](#) and [Creating a load-balancing cluster and adding components to load-balance](#).

Verifying that the Central cluster is correctly configured

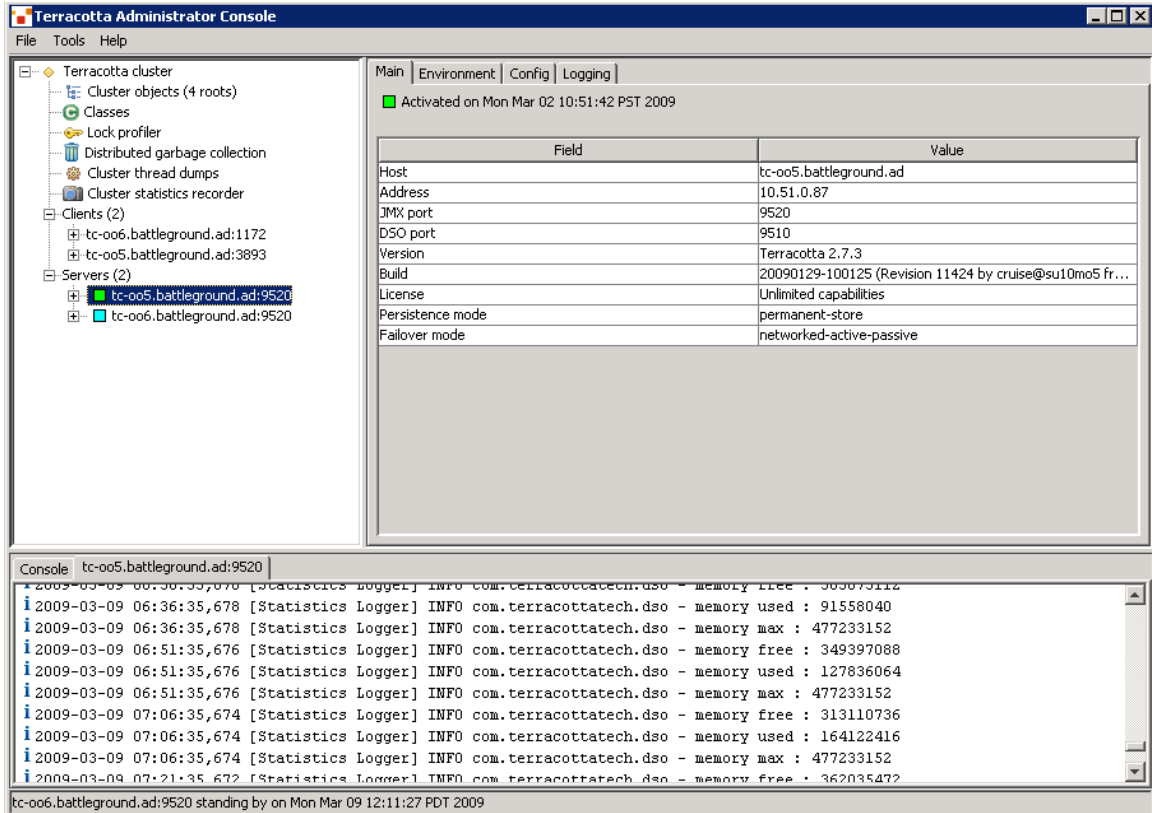
To verify that the high-availability Central cluster setup is correct, check the following:

- In the Central home directory's \conf subdirectory, in Central.properties, 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 Clustering/terracotta folder, in the 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 presented 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)").

Using 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 "Servers" and "Clients" nodes will depart from the above example.

Load-balancing for Central and RASes

You provide load-balancing for Central by installing the HP OO Load Balancer and then running the cluster-config.jar file to add Central nodes and/or standalone RAS installations to the Load Balancer configuration.

To provide load balancing for Central nodes in a cluster and/or standalone RAS installations

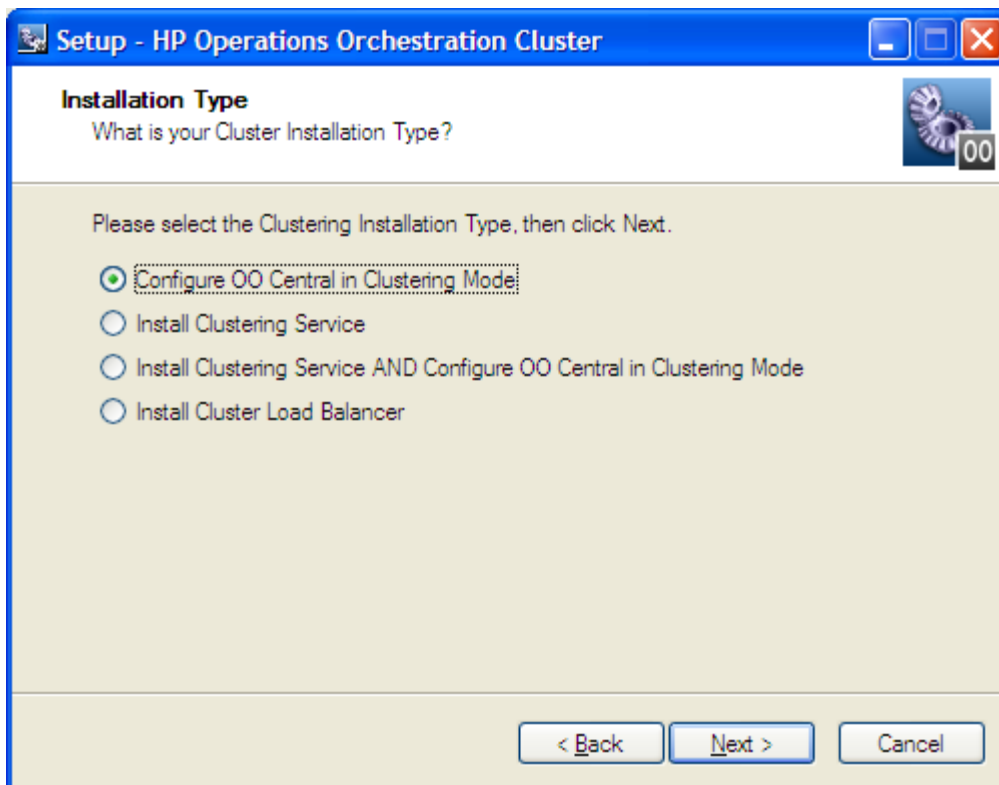
1. Install the HP OO Load Balancer, using one of the two following procedures:
 - [Installing Load Balancer on Windows](#)
 - [Installing Load Balancer on Linux](#)
2. To create the load-balancing cluster, run cluster-config.jar to manage the list of Centrals and/or RASes that you want to load-balance.

Installing Load Balancer on Windows

To install the HP OO Load Balancer on Windows

1. On the machine on which you are installing the Load Balancer, on the HP OO Installation CD, navigate to and double-click the file **ClusterInstaller-7.50.exe**.
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 then click **Next**.

The **Installation Type** page appears.



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

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

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

Installing Load Balancer on Linux

To install the HP OO Load Balancer on Linux

1. On the installation CD, in the Clustering subdirectory, locate and copy the .zip file Cluster-7.50.zip and extract it to the location where you're going to perform the installation.

Note: The following instructions reproduce the information in INSTALL.txt (contained in Cluster-7.50.zip).

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-7.50).

Using configure.sh, create or view the cluster and add or remove its nodes. The following section includes information on [creating the load-balancing cluster and adding components](#).

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

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

Creating a load-balancing cluster and adding components to load-balance

Use the following procedure in both Windows and Linux environments.

To specify which Central servers and standalone RAS installations are load-balanced

1. If you are using Windows open a command window.
2. 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.

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

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

Note: For 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

```

4. To create a cluster, type **create** and then press ENTER.
5. At the prompt **Enter the cluster name**, type the cluster name, and then press ENTER.
6. At the prompt **Enter the cluster type**, among the cluster types:
 - Type **PAS** to add a Central node.
 - OR
 - Type **RAS** to add a RAS node.

Do not specify either JRAS or NRAS.
7. At the prompt **Enter the cluster port**, type the port for the cluster (by default, this is 8443 for Central and 9004 for RAS), and then press ENTER.
A message appears, reading "Cluster <clustername> successfully created!"
8. To add a node to the cluster, type **add** and then press ENTER.
9. At the prompt **Enter the node host**, type the machine name or IP address of the node to add.
10. At the prompt **Enter the node port**, type the port that the node will use to communicate with the other nodes.
Each time you successfully add a node, the current state of the cluster, including its nodes, is displayed.
11. To save the config file, at the prompt, type **save** and then press ENTER.
The name of the cluster service on Windows systems is RSCLUSTER.
On Linux systems, the name of the load-balancer service is PASLB.sh.

Connecting Studio to the Central shared repository through the Load Balancer

Flow authors also benefit from failover clustering for Central. They access the shared Central repository through the Load Balancer. The Load Balancer thus also levels the loads imposed on Central by multiple authors accessing the shared repository simultaneously

To point Studio at the OO Load Balancer

1. In the text editor of your choice, open the Studio.properties file
The Studio.properties file is located in the OO home directory, in the \Studio\conf folder.
2. Locate the line
`dharmarepaircenter.host=<central host name>`

3. In the line, change `<central host name>` to the host name of the Load Balancer.
4. If the Load Balancer port number is different from 8443, locate the line `dharmarepaircenter.port=8443` and change 8443 to the Load Balancer's port number.
5. Save and close the file.

Troubleshooting the HP OO Cluster deployment

This section addresses one of the most common problems that may occur when deploying or expanding the cluster—the so-called “split brain” problem. 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.

The easiest way to check for a split brain situation is to:

1. Open the Terracotta Administrator Console on one of the nodes that host the Clustering Service (Terracotta Server).
2. Change directories to `<installation_location>/Clustering/terraccotta/bin`.
3. Run `admin.bat` on Windows or `admin.sh` on Linux.
4. Connect to the localhost.

A screen similar to the following should appear (if the Terracotta server on that node has correctly started).




Figure 7 - Checking for a split brain situation with the Terracotta Administrator Console

The list of servers should contain only one active node. Clicking on each server will show its state; only one should show as active in the right side panel (“Activated on....”).

If the cluster has a split brain, then the whole cluster should be shut down and the problem remediated. The clustering-specific log files (which should be looked at if a problem occurs) are located:

- 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

A common source for split brain syndrome is misconfiguration of the list of Terracotta servers in <installation_location>/Clustering/terracotta/tc-config.xml.

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

In tc-config.xml 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 <installation_location>/Clustering/terracotta/wrapper.conf 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.battleground.ad) for high availability, then its additional list should contain TC-006.battleground.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.battleground.ad" name="TC-oo5.battleground.ad">
```

...

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

A good source of information about the clustering technology, its configuration, and mode of operation is www.terracotta.org .

Index

- Central
 - and Clustering Service, arrangement, 7
- Central cluster
 - adding nodes to on Linux, 18
 - configuring, 1
 - creating, 1
 - creating on Linux, 18
 - creating on Windows, 13
 - installing, overall procedure, 6
 - upgrading, 6
 - verifying correct configuration, 20
- Central cluster nodes
 - adding on Linux, 18
 - configuring on Windows, 13
 - preparing, 8
- Central server
 - network address, 1
- Cluster
 - configuring, 1
 - creating, 1
- Cluster nodes
 - preparing, 8
- Cluster Service nodes
 - preparing, 8
- Clustering Central
 - on Linux, 16
- Clustering Service
 - and Central, arrangement, 7
- Clustering Service
 - installing on Windows, 11
- Clustering Service
 - installing on Linux, 16
- Clustering Service cluster
 - and Central cluster, creating simultaneously, 8, 13
- copyright notices, ii
- HA cluster
 - configuring, 1
 - creating, 1
 - installing, overall procedure, 6
- High Availability cluster. *See* HA cluster, *See* HA cluster
- legal notices, ii
 - copyright, ii
 - restricted rights, ii
 - trademark, ii
 - warranty, ii
- Load Balancer
 - connecting Studio to, 24
 - installing, 22
 - Installing on Linux, 23
 - installing on Windows, 22
 - upgrading, 6
- Load balancing
 - installing, overall procedure, 6
- Load-balancing cluster
 - configuring, 23
 - creating, 23
- restricted rights legend, ii
- Studio
 - connecting to Load Balancer, 24
- trademark notices, ii
- troubleshooting, 25
- warranty, ii