

HP Operations Orchestration Software

Software Version: 7.20

Installing HP Operations Orchestration

Document Release Date: July 2008

Software Release Date: July 2008



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 2008 Hewlett-Packard Development Company, L.P.

Trademark Notices

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

Updating documentation

Documentation enhancements are a continual project at Hewlett-Packard Software. You can update the documentation set at any time using the following procedure (which is also available in the HP OO readme file).

To obtain HP OO documentation

1. On The Web site <https://support1.opsware.com/support/index.php>, log in with account name and password that you received when you purchased HP OO.
2. On the **Support** tab, click the **Product Docs** subtab.
3. Under **Quick Jump**, click **Operations Orchestration** (or **Process Automation System**).
4. Under **Operations Orchestration**, click **ZIP** beside **HP OO 7.20 Full Documentation Set**.
5. Extract the files in the .zip file to the appropriate locations on your system:
 - For the tutorials to run, you must store the .swf file and the .html file in the same directory.
 - To obtain the repository that reflects the state of the flow at the start of the tutorial, unzip the file Exportof<preceding_tutorial_name>.zip.
 - To obtain the scriptlet for the tutorial that includes using scriptlets, click the scriptlet .txt file name.
 - To update your Central or Studio Help:
 - a. Under **Help Files**, click **Studio Help File Bundle** or **Central Help File Bundle**.
 - b. In the **File Download** box appears, click either **Open** or **Save**.
 - c. Extract the files to the Hewlett-Packard Software\HP OO home directory, in either the **\Central\docs\help\Central** or **\Studio\docs\help\Studio** subdirectory, overwriting the existing files.

Support

For support information, including patches, troubleshooting aids, support contract management, product manuals and more, visit one of the two following sites:

- <https://support1.opsware.com/support/index.php>
- http://www.hp.com/go/hpsoftware/DCA_support

Contents

Warranty	ii
Restricted Rights Legend	ii
Trademark Notices	ii
Support.....	iii
Overview	1
Roadmap to this guide	1
Planning the installation	1
Useful information	2
Ports of importance to HP OO components	2
Crossing firewalls.....	3
Server and Studio distribution	3
One staging Central server, one production Central server.....	3
Multiple authors	4
Multiple authors, multiple Central production servers.....	4
Extensibility	5
Creating a security certificate.....	5
Basic HP OO installation.....	5
High-level view of installing HP OO	6
High-level view of installing HP OO with clustering	7
Configuring the database and database user.....	9
Oracle 10g Enterprise: schema and schema user configuration	9
MySQL: database and database user configuration	9
SQL Server: database and database user configuration.....	9
Installing Central	11
Installing Central on a Windows operating system.....	11
Installing Central on a Linux operating system.....	15
Installing Studio	17

Further installation testing	19
Making HP OO available to users	19
Installing RAS on standalone servers	20
Installing standalone RAS on Windows.....	20
Installing standalone RAS on Linux.....	21
Testing your standalone RAS installation	22
Accessing your standalone RAS from Studio	22
Clusters and HP OO.....	22
Creating a load-balancing cluster for Central.....	24
Installing the HP OO Load Balancer.....	25
Creating and configuring a Central cluster for load-balancing	26
Creating a failover cluster for HP OO Load Balancer.....	27
Creating a Central cluster for failover and run recovery	27
Verifying that the Central failover/run-recovery cluster is correctly configured	30
Replicating the repository across the nodes of the cluster.....	30
Configuring Linux Virtual Server load balancer to use HP OO.....	31
Uninstalling HP OO programs	32
Index	33

Overview

Hewlett-Packard Software Operations Orchestration (HP OO) is a system for creating automated procedures (called *Ops flows* or *flows*) and making them available to IT personnel. Ops flows are made up of operations that each perform a specific action to maintain, diagnose, and/or repair a system.

Operations Orchestration is made up of the HP OO Central Web application, the Central Web client, and HP OO Studio:

- The Central Web application provides the business-logic tier between the HP OO Repository database and both the Central Web client and Studio.
- The Central Web client is the browser-based tool that IT personnel use to execute Ops flows.
- HP OO Studio is a standalone application in which you create new Ops flows.

Roadmap to this guide

Installing HP OO for importing, creating, and running basic Ops flows involves installing the Central Web application and Studio. The following sections of this installation guide describe these tasks in detail and provide background information:

- [Planning the installation](#)

This section discusses considerations regarding server distribution and HP OO extensibility.

Important: To run the Windows Accelerator Pack Ops flows against a remote machine, you must install the Remote Action Service on the remote machine.

- [Basic HP OO installation](#)

Basic installation includes the Central Web application and Studio.

This section includes procedures for installing Central and your Remote Action Service on Linux systems as well as on Windows.

- [Making HP OO available to users](#)

This section talks about what you need to provide to Central users and Studio authors for them to get started.

- [Clusters and HP OO](#)

Installing a load-balancing cluster and enabling failover and run recovery in a Central cluster.

For information on ongoing administrative tasks and on configuring HP OO for extended functionality, see the HP OO *Administration Guide* (AdminGuide.pdf).

Planning the installation

Planning issues include the following:

- Server and Studio distribution
- Ports for crossing firewalls
- Extensibility
- Clustering for scalability and high availability:
 - Whether to cluster Central servers for load-balancing and/or failover support
If you create a load-balancing cluster for Central servers, whether you will use HP OO Load-Balancer or a third-party load-balancer.
 - If you create a failover/run recovery cluster of Central servers, whether your Central servers can use IP multicasting (the alternative for multicasting is TCP ping).
 - Whether to cluster Central database servers
 - Whether to cluster computers on which RAS is installed
 - Whether the Central server nodes support IP multicasting
If not, then the Central nodes use TCP Ping.
- Language version
HP OO supports any language that your database and operating systems support. International language support means that with proper configuration, you can enter and obtain data in the language of your choice. To do so, all the items in the following list must be configured for the same language and character set:
 - HP OO Central: The operating system of the server on which it is installed.
 - The database that Central uses.
For information on how to configure your database to take advantage of this for MySQL, Oracle, or SQL Server databases, see [Configuring the database and database user](#).
 - HP OO Studio: The operating system of the machine on which it is installed.
 - Remote Action Service: If installed standalone, the operating system of machine on which it is installed.
 - Data sources on any machine that a flow exchanges data with. Data sources might be operating systems, databases, files, among other sources.

Note: If you install remote action services, do not install the Web server in a subdirectory of the directory that will be the HP OO home directory.

Useful information

The following information is useful for planning your installation and for various uses, such as installing or configuring HP OO components and creating URLs that launch flows.

Ports of importance to HP OO components

By default, HP OO components use the following ports for communication:

- Central: 8443
If Central servers are clustered, port 45566 is used for SSL communications between JGroups nodes

- Between HP OO components, such as Central, Scheduler, and RAS: 18443
- RAS: 9004
- Scheduler: 19443
- SMTP server: 25

Crossing firewalls

When you install RAS on the other side of a firewall from Central, you must specify a port in the firewall through which Central can communicate with the remote RAS. For a standalone RAS, this is 9004.

Also, in order to use operations that send email, your firewall must allow outbound communication over the default port used by your email server. For an SMTP server, this is port 25.

Server and Studio distribution

Factors that can influence your server distribution include whether you want to:

- Enable Ops flow authors to edit an Ops flow while users are running the previous version of the Ops flow. This requires installing Central on a staging server as well as on a production server.
- Install RAS on a standalone server. You cannot install a standalone RAS on the Central Web server; the Central installation program includes installation of RAS.

With repository publishing and updating, it is possible to have multiple authors working on Ops flows. It is recommended that in a multi-authoring installation, you install Central on a staging server for final-testing purposes as well as on the production server.

Following are some scenarios and the server and client distribution that would support them.

Important: It is best practice, particularly when you have multiple authors, to create a staging HP OO environment, from which an HP OO administrator then publishes to the production HP OO environment.

One staging Central server, one production Central server

This scenario is the minimum recommended best practice. Even if the Ops flows author and the Central user are a single person, there should be a staging server for testing and a production server that the repository is published to after testing on the staging server. Studio can be installed either on a third computer, on the staging server, or on the production server.




Figure 1 - Minimum recommended configuration

The Central staging server and the Central production server each have Central installed. Studio can be installed on either server. Whether Studio is installed on the Central production or staging server, the author publishes to the repository that the staging server's installation of Central uses.

Multiple authors

Multiple authors can exchange work by publishing and updating from the Central staging server. When either author publishes to the staging server, his or her work is copied there. When either author updates from the staging server, his or her local repository is updated with any work that the other author has published to the staging server.




Figure 2 - Server distribution for multiple authors

Multiple authors, multiple Central production servers

If you cluster two or more Central production servers, your server distribution would look like the following, if you had two authors. The authors publish to and update from the staging server as described in the preceding scenario. The difference here is that the staging server's Central repository is published to either or both of the clustered Central production servers. The relationship between the clustered production servers depends on the clustering software that is used.




Figure 3 - Server distribution for multiple authors and multiple production servers

Extensibility

Using basic HP OO operations, you can carry out Ops flows that are enabled through the Web service on which runs. Using RAS, you can extend the ability of HP OO to operate in additional contexts:

- On machines that are remote from the Central Web server. Such machines can be on the other side of a firewall from the Central Web server or anywhere on an intranet or the Internet.
- Using Web services other than that on which the Central Web server was installed.
- Integrated with actions carried out on other applications, programmed using the target application's API.

Creating a security certificate

For installations of Central that communicate using the HTTPS protocol, Web browsers will show security violation errors or messages unless you create a valid security certificate for delivering the Central Web pages. If you see such a browser warning, it is because HP OO includes, by default, an unsigned certificate that serves as a placeholder for a valid customer-obtained certificate. If you choose not to create a security certificate, you can safely ignore the warning.

Basic HP OO installation

Installing Central and Studio provides everything you need to create and run Ops flows. The Central installation includes automatic installation of the Remote Action Service (RAS), which enables you to run Ops flows that operate remotely from HP OO or interact with other programs.

High-level view of installing HP OO

The following list of tasks assumes that you want to have the following, but without load-balancing or failover clustering support:

- Multiple authors publishing to a Central installation in a testing (or staging) environment.
- A remote RAS installation.

If you are only installing Central on a single server, you only need to complete steps 1 and 3.

You would install and configure HP OO components in the following order:

1. Make sure that the machines on which you install the HP OO components meet hardware and software requirements described in:
 - The Hewlett-Packard Software HP OO System Requirements document (SystemReqs.pdf)
 - Requirements for all the required software programs.
2. Create a testing/staging HP OO environment, which we will assume has the following configuration:
 - Database server
 - Central server
 - Optional: installation of RAS on its own computer

The high-level steps to create this testing or staging environment are:

- a. Install the database management system on the database server.
 - b. Create a database for Central and configure the database and database user as necessary.
For information, see [Configuring the database and database user](#).
 - c. Install Central.
For information on and procedures for installing on either a Windows or a Linux system, see [Installing Central](#).
 - a. Install Studio.
For information, see [Installing Studio](#).
 - b. If necessary, install standalone RAS.
For information on and procedures for installing on either a Windows or a Linux system, see [Installing RAS on standalone servers](#).
3. Create the production HP OO environment, which we will assume has the following configuration:
 - Database server (clustered, thus more than one)
 - Central server (clustered, thus more than one)
 - Optional: installation of RAS on a remote cluster of computers

The high-level steps to create this configuration are:

- a. Install the database management system on the database servers.
- b. Create a database for Central and configure the database and database user as necessary.
For information, see [Configuring the database and database user](#).

- c. Following your database vendor's recommendations, create a failover and/or load-balancing cluster for the database servers.
- d. Install Central on the Central servers that will be nodes in the Central cluster. For information on and procedures for installing on either a Windows or a Linux system, see [Installing Central](#).
- e. Install Studio on each author's machine, referencing the Central installation. For information, see [Installing Studio](#).
- f. If necessary, install standalone RAS (clustered, thus more than one). For information on and procedures for installing on either a Windows or a Linux system, see [Installing RAS on standalone servers](#).

High-level view of installing HP OO with clustering

Suppose that you want to have:

- Multiple authors publishing to a Central installation in a testing (or staging) environment.
- Load-balancing and failover capability for the Central servers in your production environment.
- Load-balancing support for a remote RAS installation.

You would install and configure HP OO components in the following order:

1. Make sure that the machines on which you install the HP OO components meet hardware and software requirements described in:
 - The Hewlett-Packard Software HP OO System Requirements document (SystemReqs.pdf)
 - Requirements for all the required software programs.
2. Create the following testing/staging HP OO environment, which we will assume does not need load-balancing or failover support:
 - Database server
 - Central server
 - Optional: installation of RAS on its own computer

The high-level steps to create this environment are:

- a. Install the database management system on the database server.
- b. Create a database for Central and configure the database and database user as necessary. For information, see [Configuring the database and database user](#).
- c. Install Central. For information on and procedures for installing on either a Windows or a Linux system, see [Installing Central](#).
- d. Install Studio. For information, see [Installing Studio](#).
- e. If necessary, install standalone RAS. For information on and procedures for installing on either a Windows or a Linux system, see [Installing RAS on standalone servers](#).

3. Create the following production HP OO environment, which we will assume needs load-balancing or failover support:

- Database server (clustered, thus more than one)
- Central server (clustered, thus more than one)
- Optional: installation of RAS on a remote cluster of computers

The high-level steps to create this configuration are:

- a. Install the database management system on the database servers.
- b. Create a database for Central and configure the database and database user as necessary.
For information, see [Configuring the database and database user](#).
- c. Following your database vendor's recommendations, create a failover and/or load-balancing cluster for the database servers.
- d. Install Central on the Central servers that will be nodes in the Central cluster.
For information on and procedures for installing on either a Windows or a Linux system, see [Installing Central](#).
- e. Install HP OO Load Balancer (probably on one of the Central servers).
For information on and procedures for installing on either a Windows or a Linux system, see [Installing the HP OO Load Balancer](#).
- f. Using the Load Balancer, create and configure the Central load-balancing cluster.
For information, see [Creating and configuring a Central cluster for load-balancing](#).
- g. By configuring the Central.properties file on each Central server, create the Central failover/run recovery cluster.
For information, see [Creating a Central cluster for failover and run recovery](#).
- h. Create a failover cluster for HP OO Load Balancer.
For information, see [Creating a failover cluster for HP OO Load Balancer](#).
- i. Install Studio on each author's machine, referencing the Central installation.
For information, see [Installing Studio](#).
- j. If necessary, install standalone RAS (clustered, thus more than one).
For information on and procedures for installing on either a Windows or a Linux system, see [Installing RAS on standalone servers](#).
- k. Using the Load Balancer, create and configure load-balancing clusters for the RAS installations.
The method is the same as for creating a Central load-balancing cluster.
- l. Replicate the public repository across the Central cluster nodes.
For information, see [Replicating the repository across the nodes of the cluster](#).

Configuring the database and database user

You must create and configure a database and database user before installing Central because Central does not create the database, but the schema for the database.

Important: The schema and user must always be dedicated to HP OO. Do not adapt the schema to support any other applications. If you do, you can lose data in the other applications when you install HP OO.

The Central database size and collation depend on your business needs.

As you install or configure your database management system and create a database and database login and user, make sure that among your tasks you complete the tasks in one of the following sections, depending on which database management system you use.

Oracle 10g Enterprise: schema and schema user configuration

Multiple language support

To enter data in Central and obtain data from the schema in a language other than English

- Configure the database server with UTF-8.

Privileges for the database user

When creating a user for the schema, grant CONNECT and RESOURCE privileges to the user. The Central installer makes all the necessary changes to the schema that Central will use.

MySQL: database and database user configuration

Multiple language support

To enter data in Central and obtain data from the database in a language other than English

- Configure the database server with UTF-8.

Privileges for the database user

When you create a database and a user for the database, grant all privileges for that database to the database user.

SQL Server: database and database user configuration

Multiple language support

To enter data in Central and obtain data from the database in a language for which SQL Server supports localization, configure the database with the appropriate language collation.

Following are the required collations in SQL Server for the listed languages:

SQL Server version	Language	Required collation
SQL Server 2005		
	Japanese	Japanese_Unicode_CS_AS
	German	German_PhoneBook_CS_AS & Latin1_General_CS_AS
	Korean	Korean_90_CS_AS
SQL Server 2000 SP3a and above		
	Japanese	Japanese_Unicode_CS_AS
	German	German_PhoneBook_CS_AS & Latin1_General_CS_AS
	Korean	Korean_Wansung_Unicode_CS_AS

For example, to enter Japanese characters in Central and receive data in Japanese from the database, configure the database with an appropriate Japanese collation.

[Database information that the person who installs Central needs](#)

If another person installs Central, you must provide him or her with the following information about the database and the database server:

- Database name
- Server location and name or IP address
- Port number
- Database user name and password

SQL Server 2000, SP 3a or later

- When you create the login for the new database user, among the new login's properties:
 - Specify either kind of authentication.
 - Select the database that you have created as the default database for this login.

Note: If this is not acceptable for your organization, the database administrator can modify the .ddl file appropriately. The .ddl file is located in the HP OO home directory, in the \Central\db subdirectory.
- When creating the database user:
 - For the default schema, specify "dbo".
 - For role membership, select the db_owner database role.

SQL Server 2005

1. When configuring SQL Server configuration, enable TCP/IP as a network protocol.
 2. Enable TCP/IP as a network protocol for SQL Server.
 - a. Within the SQL Server Configuration Manager program, expand **SQL Server 2005 Network Configuration**.
 - b. Select **Protocols for MSSQLSERVER**, right-click **TCP/IP**, and in the context menu that appears, click **Enable**.
 3. When creating the database login:
 - Specify either kind of authentication.
 - Select the database that you have created as the default database for this login.

Note: If this is not acceptable for your organization, the database administrator can modify the .ddl file appropriately. The .ddl file is located in the HP OO home directory, in the \Central\db subdirectory.
 - In **Server Roles**, select all the roles.
 4. When creating the database user:
 - For the default schema, specify "dbo".
 - For role membership, select the db_owner database role.
- Next, you install the Central Web application and Studio.

Installing Central

You will find files for installing Central on either a Windows or a Linux system. The procedure for installing on a Linux system.

In the course of the installation, the program:

- Checks connectivity with the database, using the database connection information that has been provided.
- Configures the database schema that has been created.
- Installs the Central server files.
- Installs RAS on the Central server.

Note: By default, Central is configured to communicate with standalone installations of RAS over port 9004, respectively (which the standalone RAS installation configures on the remote machines for communication with Central).

At any point in the following procedure, you can stop the installation by clicking **Quit**. The installation stops gracefully.

Important: If you are installing using a MySQL database, the installation requires that you have the Java database connectivity (JDBC) driver file FOR MySQL.

Installing Central on a Windows operating system

To install Central on a Windows operating system

1. Before starting the installation, make sure that you are connected to the database that you're using for Central.

2. If Central will use a MySQL database, obtain the MySQL JDBC driver file `Mysql-connector-java-<versionnumber>.jar` and copy it to a location of your choosing.
3. Navigate to and double-click `CentralInstaller-<version_number>.exe`.
4. On the **Welcome** page of the Central Setup Wizard, click **Next**.
5. On the **License Agreement** page, read the agreement, select **I accept the agreement**, and then click **Next**.

The **Select Destination Location** page appears, unless you are reinstalling the Web application or have already installed Studio on this machine.

If you are have already installed Central on this machine, the **Select Destination Location** page does not appear. Rather, Central is automatically installed in the Central subfolder of the HP OO home directory (in the Studio subfolder of which Studio is installed). If this is the case, skip the following step 6 and resume this procedure where the **Network Settings** page appears.

6. On the **Select Destination Location** page, click **Next** to accept the default installation path.

OR

Click **Browse**, specify a different location where you want the Web application installed, and then click **Next**.

The **Network Settings** page appears, on which you configure the ports and IP address that Central uses.

Setup - HP Operations Orchestration Central

Network Settings

Please provide the following information and then press Next

Please provide the HTTP and HTTPS ports to be used by HP Operations Orchestration Central. All sensitive data will use HTTPS protocol (encrypted network traffic).

Central HTTP Port Number (usually 80 or 8080):
8080

Central HTTPS Port Number (usually 443 or 8443):
8443

Broadcast address (e.g. 192.168.1.255):
192.168.1.255

< Back Next > Cancel

Figure 4 – Network Settings page

Communications between the Central Web application, Studio, and the Central Web client are over secure connections, using the HTTPS protocol. The HTTP port

makes it possible for client users to type their customary HTTP protocol in their browser address box and be redirected to the secure connection.

The broadcast address is the address that RASs will use to broadcast their availability to Studio and Central. (For more on RASs and checking their availability, see the HP OO Concepts Guide (ConceptsGuide.pdf) and Help for Studio.

7. Either accept the default values or type different port numbers in the text boxes.
8. Record the HTTPS port number in a separate document.

Note: The installation program for Studio requires the HTTPS port number, so make sure it is available for installations of Studio.

9. Either accept the default broadcast IP address or type a different one, and then click **Next**.

The first **Database Information** page appears.

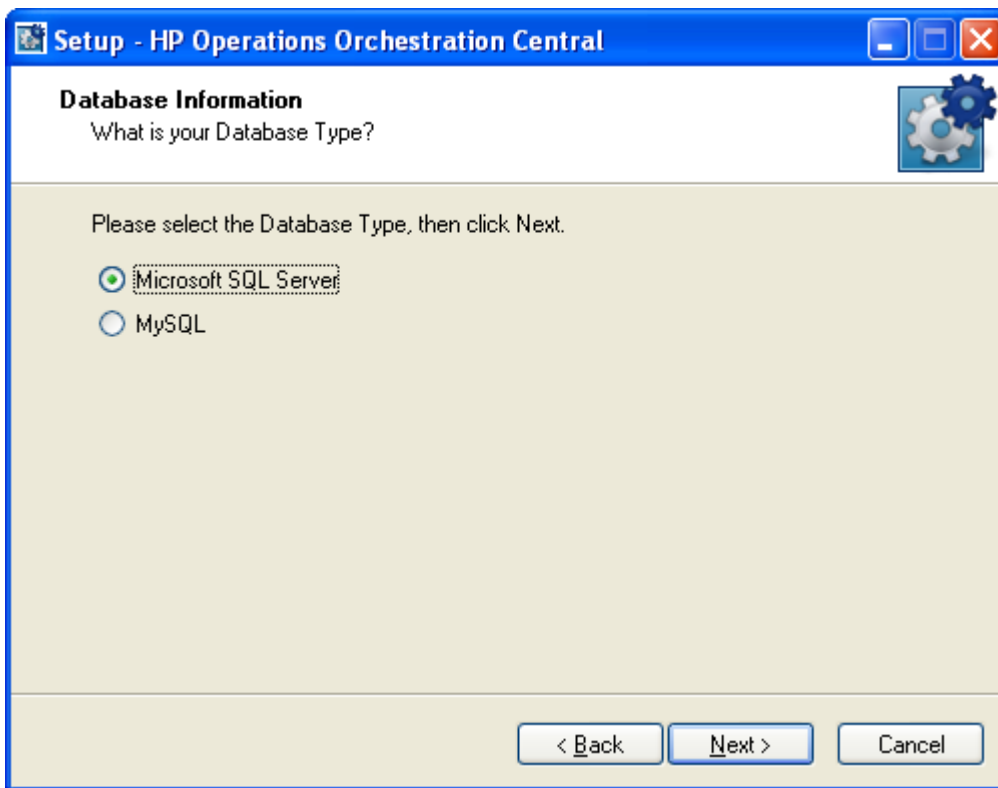


Figure 5 - Database Information page

10. Select the type of database that Central will access, and then click **Next**.

The second **Database Information** page appears.

Warning: If you choose to create a new database with the name of an existing database, all the data in the previously installed database is overwritten.

11. Fill in the information as your configuration dictates.

Use the following formats. In these formats, angle brackets (<>) contain strings that you replace with your particulars; square brackets ([]) indicate optional material; and the pipe character (|) separates two alternatives.

- **DB Host or IP Address**

<host|ip_address>

- **DB Port Number**

<port>

Be sure to record the port number that your database instance uses; it will be needed for installing Studio.

- **DB Name**

<db_name>[;instance=<instance_name>]

The instance name is necessary only if the database is an instance.

- **DB User Name**

[<domain_name>\]<db_user_name>

The domain name is necessary only if the database authentication is Windows, using a domain account.

You can get this information from your database administrator.

- **Password and Confirm Password**

The database name (**DB Name**) and password are the name and password that the Central Web application uses to access the database.

12. To create a new database, select the **Create (or recreate, if it exists) the database schema** checkbox, and then click **Next**.

The Installation program tests the database connection using the information that you provided. The **Test Database Connection** progress page appears and shows progress of the database connection test.

13. If the database connection test fails, return to the **Database Information** page and correct the information.

After the database connection test succeeds, if you have chosen to create or recreate the database, the **Create HP OO Administrator Account** page appears. The administrator account, which you will use to log in to Central, is always named "admin."

OR

If you are reinstalling Central and have not chosen to create or recreate the database, then this page does not appear. (Once you have created the admin account and its password, you cannot change the password.) Instead, the installation program goes directly to the **Ready to Install** page. In this case, skip the following step.

14. In the **Password** and **Confirm Password** text boxes, create a password for the admin account, and then click **Next**.

The **Ready to Install** page appears, displaying the information that you specified in the rest of the setup.

15. To proceed, click **Install**.

The installation program begins installing, tracking its progress on the **Installing** page.

16. When the installation completes, click **Finish**.

Next, you install Studio. For the Studio installation procedure, see [Installing Studio](#).

Installing Central on a Linux operating system

You can have as many HP OO Central installations as you wish, all running on the same host; these installations can be of the same version or different versions. When installing multiple copies of HP OO on one host, edit the install.config files for each installation to make sure that there are no port conflicts and no database conflicts among the copies of HP OO.

There are two versions of this installation package for each database management system, one for the 64-bit version of Linux and one for the 32-bit version:

- On a 64-bit version of Linux, you should probably be using the 64-bit version of the installation .zip file (PASCentral-7.10-linux64-<database system>-* .zip).
- On a 32-bit version of Linux, be sure to use the 32-bit version of the installation .zip file (PASCentral-7.10-linux32-<database system>-* .zip).

This package includes Java Runtime Environment (JRE), version 1.6.

To install Central on a Linux operating system

1. From the installation files, unzip the appropriate .zip file **Central-7.20-linux<32|64>-<databasetype>.zip**

Where:

<32|64> represents the Central installation for 32- or 64-bit systems.

<dbasetype> stands for the type of database management system that you will use, either MySQL or Oracle.

2. Upgrade from 7.0.

For information on upgrading, see the HP OO 7.10 Upgrade Guide.

3. Edit the install.config file and change parameters to fit your environment.

Note: If the database used is an Oracle RAC configuration, the settings for DB_HOST and DB_PORT are not used. This is because a cluster usually involves two or more hosts with possibly different ports. Thus, the settings for ORACLE_RAC_NODES and ORACLE_RAC_ONS_CONFIG must be uncommented and populated accordingly. If by mistake both DB_HOST/DB_PORT and ORACLE_RAC_NODES/ORACLE_RAC_ONS_CONFIG are populated, the RAC settings will take precedence, whereas the other settings will be ignored.

4. Run the script ./configure.sh in this directory.

This script will apply the configuration parameters to the appropriate files and will optionally create or recreate the database schema. It also produces the .iconcluderc file in the current directory.

The following log files are written by configure.sh:

- ./configure.log
- ./recreate-db.log

5. Source the file .iconcluderc and verify that the environment variable \$ICONCLUDE_HOME is properly set.

6. Start the Central service:

```
$> $ICONCLUDE_HOME/bin/Central.sh start
```

Note: the options for all the startup scripts are:

- console
Runs the program in console mode (non-daemon, has controlling tty).
- start
Starts as daemon.
- stop
Stops program.
- restart
Stops, then starts as daemon.
- status
Shows status and PID.
- dump
Sends a SIGQUIT to the java process, forcing it to do a thread dump.

7. Start the RAS service (optional):

```
$> $ICONCLUDE_HOME/bin/JRAS.sh start
```

The options for Central.sh can also be used with this command.

8. If you wish to increase the capacity of the system:

- a. Increase the fd limit for the user under which Central runs, by adding the following two lines to the file /etc/security/limits.conf:

```
myuser soft nofile 65536
myuser hard nofile 65536
```

- b. Log in as "myuser" and check the value by issuing the `ulimit -n` command.

The following are useful log files to monitor:

- \$ICONCLUDE_HOME/Central/logs/Central_wrapper.log
- \$ICONCLUDE_HOME/RAS/Java/Default/webapp/logs/wrapper.log

Other useful locations are:

- \$ICONCLUDE_HOME/bin has symbolic links to startup files for Central and RAS.
- \$ICONCLUDE_HOME/conf has symbolic links to wrapper.conf files for Central and RAS.

Next, you install Studio. For the Studio installation procedure, see [Installing Studio](#).

Installing Studio

To install Studio, you must be logged in with an account that is either a Power User or a local Administrator.

Important: The installation requires a port number for communications using the HTTPS protocol. This port number was specified during installation of the Central Web application. Before starting installation of Studio, obtain this port number from the person who installed Central.

To install Studio

1. To start the Studio Setup Wizard, navigate to and double-click StudioInstaller-`<version_number>.exe`.
2. On the **Welcome** page, click **Next**.
3. On the **License Agreement** page, accept the terms of the license agreement, and then click **Next**.

The **Select Destination Location** page appears, unless you are reinstalling the Web application or have already installed Studio on this machine.

If you have already installed Central on this machine, the **Select Destination Location** page does not appear. Rather, Studio is automatically installed in the Studio subfolder of the HP OO home directory (in the Central subfolder of which Central is installed). If this is the case, skip the following step 4 and resume this procedure where the **HP Operations Orchestration Information** page appears.

4. On the **Select Destination Location** page, type the location where you want the Studio files installed.

OR

Click **Browse**, specify a different location where you want Studio installed, and then click **Next**.

The **HP Operations Orchestration Information** page appears.

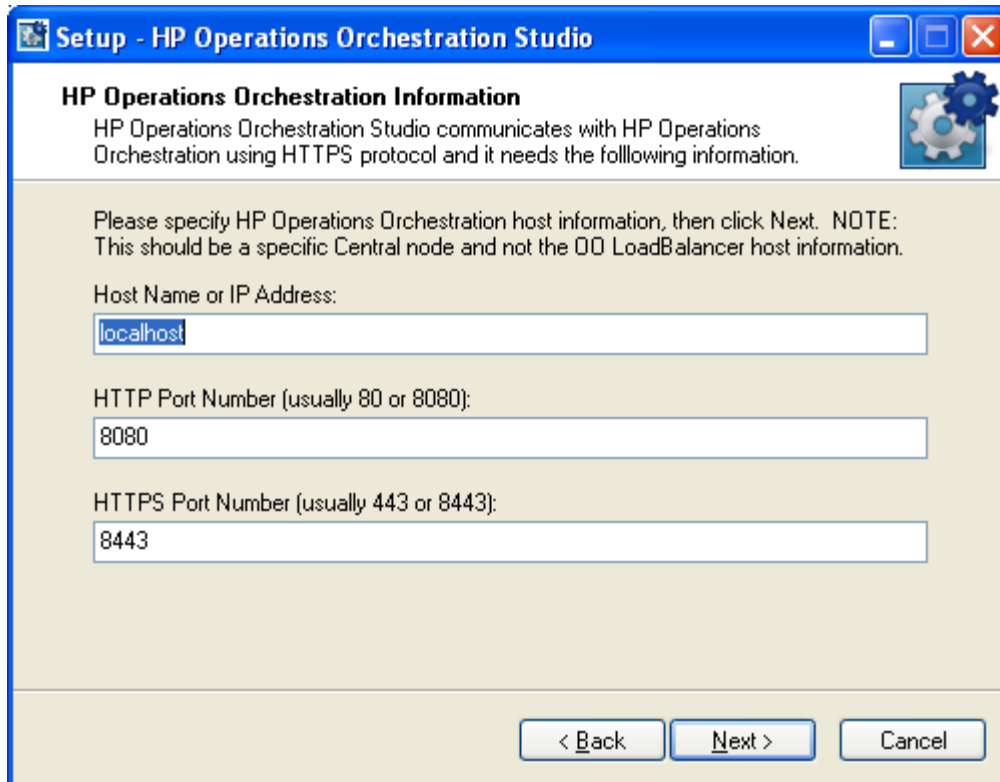


Figure 6 – HP OO Information page

Important: If you install the HP OO Load Balancer, then in the following step, do not enter the URL for the Load Balancer in the **Host Name or IP Address** box. Enter the information for the actual Central server. Studio cannot access the Central server through the Load Balancer URL.

5. To specify the server on which the Web application is installed, type the server name in the text box beside **Host Name or IP Address**.

OR

If you are installing the Studio on the same machine on which the Web application is installed, leave **localhost** in the text box.

6. Make sure the HTTPS port number matches the HTTPS port number that the administrator specified when installing Central, and then click **Next**.

The **Select Additional Tasks** page appears.

7. To create a desktop icon for Studio, select the **Create a desktop icon** checkbox, and then click **Next**.

The **Ready to Install** page appears, displaying the location where Central will be installed and other information on the choices you have made.

8. To proceed, click **Install**.

The Setup Wizard tracks progress on the **Installing** page.

The installation completes, and the **Completing the...Setup** page appears.

9. To test the success of the installation, select the **Launch HP OO Studio** box, and then click **Finish**.

10. Switch to Studio and log in with the username and password that were created in the Central installation.

If you did not install Central, get the username and password from the person who did.

Further installation testing

Testing the installation involves the following high-level steps. For information on performing the necessary tasks in Studio and Central, see Help for Studio and Help for Central.

1. After logging in to Studio, in the Library pane, open the **Library** folder.
2. In your Web browser, access and log in to Central.
3. Click the **Ops Flows** tab, and open some Ops flows in the Library.

Making HP OO available to users

For HP OO users to log in to Central or start Studio on their desktop:

- The Central Web server must be started.
- Central users need the URL and login information for the Central Web site.
- Studio users (authors) need the login information and, if Studio is not already installed on their machine, the installation files.

Before you provide them with the information and files, HP OO, make sure that client machines meet hardware and software requirements for Central or Studio, as described in the *HP OO System Requirements* document, *SystemReqs.pdf*, and requirements for any required software.

To deploy Central and Studio to users

1. To start the Web server that contains the Central Web application, make sure that the RSCentral Windows service is started.
2. Send each Central user the following:

- The URL for the Central Web site.

This URL has the format:

`http://<servername>:<portnumber>/PAS/`

where <servername> is the name of the server on which the application is configured and <portnumber> is the port number that you specified when creating the application. By default, this port number is 8080 for HTTP or 8443 for HTTPS.

- The administrator login credentials (user name and password) that you created when installing the Central Web application.

Each user must log in with these credentials until you map his or her user account to the HP OO ADMINISTRATOR role.

3. Send each Studio user (Ops flow author) the following:

- If you have not installed the author's copy of Studio, StudioInstaller.exe

- The administrator login credentials (user name and password) that you created when installing the Central Web application
Each user must log in with these credentials until you map his or her user account to the HP OO ADMINISTRATOR role.

Installing RAS on standalone servers

You can install an external Remote Action Service (RAS) on a computer that is remote from the Central server. The RAS is enabled for running operations that need either the Java Runtime Environment (version 1.6) or the .NET framework.

Note: By default, the standalone RAS installation configures it to communicate over port 9004, which is also configured on the Central server for communicating with RAS.

Because RAS is installed as part of the Central install, this installation can only be done on machines on which Central has not been installed. This installation includes the installation of any HP OO content included in this release.

Installing standalone RAS on Windows

To install the RAS server

1. On the Hewlett-Packard Software HP OO CD, locate RASInstaller-
<versionnumber>.exe (the RAS installation executable file) and copy it to the machine on which you're going to install it.
The RAS installation .exe may be in the JRAS folder, even though the RAS that it installs is enabled for both Java or .NET.
2. Copy the installation .exe to the machine on which you are going to install it.
3. To start the installation .exe, double-click the file.
If Central is installed on the machine on which you're running the installation program, you are warned that you must uninstall Central on the machine before installing RAS. Either click **Yes** or click **No** and move the RAS installation .exe to a machine on which Central is not installed, then restart the installation program.
The HP OO RAS Setup Wizard starts with the **Welcome** page.
4. Click **Next**.
5. On the **License Agreement** page, read the agreement, click **I accept the agreement**, and then click **Next**.
The **Network Settings** page appears.
6. In the **Broadcast Address** box, type the IP address that HP OO uses to communicate with other machines, and then click **Next**.
If you are installing on a machine on which folders remain from a previous HP OO installation, the Setup program informs you that it will install RAS using that folder structure, and that you will not be prompted for a destination location.
The Select Start Menu Folder page prompts you to create or select for the folder to hold the RAS program's shortcuts.

7. To accept the default folder, click **Next**.
OR
To choose a new folder, click **Browse**, navigate to and select the folder, click OK to return to this page, and then click **Next**.
If there are not folders left over from an earlier HP OO installation, the **Select Destination Location** page appears.
8. If the **Select Destination Location** page appears at this point, either accept the default installation path for RAS or click **Browse** and specify a different location where you want RAS installed—then click **Next**.
The **Ready to Install** page appears, displaying the location where RAS will be installed and other information on the choices you have made.
9. To proceed, click **Install**.
The installation program tracks progress on the installation progress page.
10. When the installation completes, click **Finish**.

Installing standalone RAS on Linux

To install a standalone RAS on a Linux operating system

1. Edit the file `install.config` and change parameters to fit your environment
Note: This package includes Java Runtime Environment (JRE), version 1.6.
2. Run the script `./configure.sh` in this directory.
This script will:
 - Apply the configuration parameters to the appropriate files.
 - Produce a file called `.iconcluderc` in the current directory.
 - Write the `./configure.log` file.
3. Source the file `.iconcluderc` and verify that the environment variable `$ICONCLUDE_JRAS_HOME` is properly set.
4. Start the server:

```
$> $ICONCLUDE_JRAS_HOME/bin/JRAS.sh start
```

Note: The options for all the startup scripts are:
 - `console`
Runs the program in console mode (non-daemon, has controlling tty).
 - `start`
Starts as daemon.
 - `stop`
Stops program.
 - `restart`
Stops, then starts as daemon.
 - `status`
Show status and PID.
 - `dump`

Sends a SIGQUIT to the java process, forcing it to do a thread dump.

5. If you wish to increase the capacity of the system:
 - a. Increase the fd limit for the user under which JRAS runs, by adding the following two lines to the file `/etc/security/limits.conf`:

```
myuser soft nofile 65536
myuser hard nofile 65536
```
 - b. Log in as "iconclude" and check the value by issuing the `ulimit -n` command.

The following are useful log files to monitor:

- `$ICONCLUDE_JRAS_HOME/jetty/extra/linux/iConclude.log`
- `$ICONCLUDE_JRAS_HOME/jetty/logs/wrapper.log`
- `$ICONCLUDE_JRAS_HOME/RAS/Java/Default/webapp/logs/wrapper.log`

Other useful locations are:

- `$ICONCLUDE_JRAS_HOME/SDK` contains files supporting RAS development.
- `$ICONCLUDE_JRAS_HOME/bin` has symbolic links to the startup file
- `$ICONCLUDE_JRAS_HOME/conf` has symbolic links to the wrapper.conf file

Testing your standalone RAS installation

To test the RAS installation

- Using a Web browser, access the following URL to confirm that the Web service is running.

`http://<Hostname>:<Portnumber>/RAS/services/RCAgentService?wsdl`

where

- `<Hostname>` is the name of the Central Web application server.
- `<Portnumber>` is the name of the default port used for RAS by the Central server.

By default in the Central installation, this port number is 9004.

If the Web service is running, a WSDL/XML document appears.

Accessing your standalone RAS from Studio

To obtain access to the RAS interface within Studio

- Open Studio and configure a RAS reference.
For information on configuring a RAS reference, see Help for Studio.

Clusters and HP OO

If you run large volumes of HP OO flows and if high availability is a requirement for your HP OO installation, you can create a cluster that provides your HP OO

configuration with load-balancing and failover clustering support in any combination of the following:

- A load-balancing cluster for Central servers and any standalone RAS installations
You can use HP OO Load Balancer to create, configure, and run such a load-balancing cluster.
- Failover support for the HP OO Load Balancer
You can use the clustering software of your choice to provide this failover support.
- Failover support of Central servers and automatic run recovery (which means if one of the cluster members fails, one of the other nodes will take over and resume the failed node's runs).
You can configure Central to create the failover/run recovery support for the Central servers.
- A load-balancing and/or failover cluster for the Central database servers, which contains the Central nodes' run histories, schedules, and system configuration, such as enabling of user authentication providers.
You can use any clustering software to cluster the Central database servers.

Each Central node must be its own dedicated machine, and use its own repository (a hierarchical folder structure of XML files). HP OO provides flows or operations with which you can synchronize repositories and iActions between all the nodes of the application cluster.

To combine multiple-authoring with a Central cluster, install Central in a staging environment to which the authors will publish and from which they will update. This staging instance of Central should be the only source for publishing flows, iActions, and other HP OO objects to the production environment. You can do this and maintain consistency between the nodes' repositories by one of the following:

- Publish from the staging Central server to each of the nodes in the production Central cluster.
- Run a flow that publishes from the staging Central repository to the repositories on the nodes.

Each Central node, by default, can run 100 concurrent connections either through browsers or automatically initiated runs. If your cluster is handling more requests than that, any remaining boxes should be able to handle your peak load during a flood.

Following is a schematic of what a clustered HP OO installation might look like. This drawing is based on the multi-author/staging Central/clustered production Central server configuration described and sketched above, in *Multiple authors, multiple Central production servers*. Note, however, that this schematic assumes that one of the Remote Action Services (RAS) is installed on a standalone host.

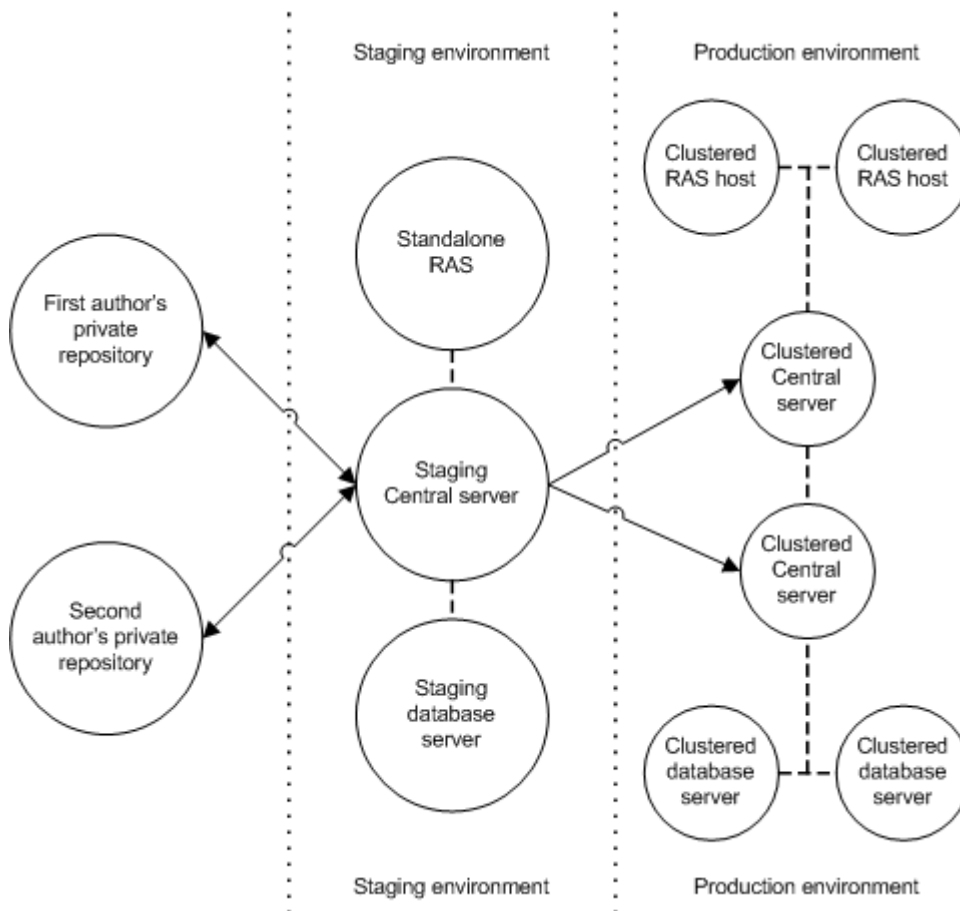


Figure 7 - HP OO configuration with staging environment and clustered production environment

Creating a load-balancing cluster for Central

Creating a load-balancing cluster for Central involves the following steps:

1. Installing the HP OO Load Balancer
2. Creating and configuring a Central cluster for load-balancing
3. Creating a failover cluster for the HP OO Load Balancer
4. Creating and configuring a Central cluster for failover/run recovery
5. Replicating the repository across the nodes of the Central cluster

If you were to install and configure load-balancing and failover clustering in a minimal fashion on two servers, the installed and configured clusters and HP OO programs might look like the following:

- Server 1
 - Windows/Linux Cluster1 in active/passive mode
 - Load Balancer 1
 - Central 1
 - RAS 1
- Server 2

- Windows/Linux Cluster1 in active/passive mode
- Load Balancer 2
- Central 2
- RAS 2

Server 1 Load Balancer would handle all HTTP requests and route them to Central 1 or Central 2, RAS 1 or RAS 2. Server 2 Load Balancer would be inactive.

If Server 1 goes down, then Server 2 Load Balancer would become active and handle all HTTP requests, sending them to Central 2, and RAS 2, as long as Server 1 is down.

Installing the HP OO Load Balancer

The HP OO Load Balancer is a Java command-line utility whose installation varies according to whether you install it in a Windows or a Linux environment. We'll look at a Windows installation first.

To install the HP OO Load Balancer for Central on a Windows operating system

1. To start the Clustering Setup Wizard, navigate to and double-click ClusterInstaller.exe.
2. On the **Welcome** page, click **Next**.
3. On the **License Agreement** page, accept the terms of the license agreement, and then click **Next**.
4. 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 HP OO Cluster to be installed.
 - Click **Browse**, specify a different location.

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

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

To install the HP OO Load Balancer for Central on a Linux operating system

1. On the installation CD, in the Clustering subdirectory, locate and copy the .zip file Cluster-7.20.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-7.20.zip).

2. Run configure.sh script.
3. Use the config-clustering.jar tool to create the cluster.
4. Run the PASLB.sh start command to start the service.

Creating and configuring a Central cluster for load-balancing

You both create and configure the Central load-balancing cluster with the HP OO Load-Balancer.

To configure a cluster

1. Open a command window.
2. Navigate to the location where you installed the HP OO Load Balancer.
3. Open cluster-config.jar, using the following:

```
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**, type the cluster type (Central and RAS, and then press ENTER.
7. At the prompt **Enter the cluster port**, type the port for the cluster (by default, this is 8443) , and then press ENTER.

A message appears, reading "Cluster <clustername> successfully created!"

```
Enter a Command: create
Enter the cluster name: mycluster
Enter the cluster type (PAS, JRAS, NRAS): PAS
Enter the cluster port: 8443
Cluster mycluster successfully added!
```

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 one of the nodes.
10. At the prompt **Enter the node port**, type the port that the node will use to communicate with the other nodes.

This can be the same port number as the port that you specified for the cluster.

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: hamlet.battleground.ad
Enter the node port: 8443
Node hamlet.battleground.ad successfully added!
Name: mycluster
Host: P4S
Port: 8443
Nodes:
  Port      Host
  ----      -
1      8443     127.0.0.1
2      8443     hamlet.battleground.ad
```

11. To save the config file, at the prompt, type **save** and then press ENTER.

Creating a failover cluster for HP OO Load Balancer

Creating an active-passive failover cluster for the HP OO Load Balancer provides high availability for the Load Balancer.

To create a failover cluster for HP OO Load Balancer

1. Install HP OO Load Balancer on each node of the cluster.
For instructions on installing HP OO Load Balancer, see [Installing the HP OO Load Balancer](#), above.
2. Create the load-balancing cluster in the HP OO Load Balancer installation on one node, as described in [Creating and configuring a Central cluster for load-balancing](#).
3. Copy the HP OO.conf file from the node on which you created the load-balancing cluster to the other nodes.
4. Create the failover cluster for HP OO Load Balancer, using the clustering technology of your choice.
See the documentation for your clustering technology for instructions on creating the cluster.

Creating a Central cluster for failover and run recovery

The settings for creating and configuring a failover cluster for Central are specified from with Central, on the **Administration** tab.

You can define one or more clusters of Central servers, each cluster providing failover capability and automatic recovery of runs that were abandoned as a result of the failure of one of the nodes. When one of the cluster nodes fails, its active runs are paused and moved into an IDLE state in Central. Then Central on the alternate node resumes the run.

To enable failover and automatic run-recovery, you configure each member, or node, of the Central cluster, on the Central Administration tab.

Among the configurations, you will specify one of the two following protocols for the cluster nodes to communicate with each other:

- IP multicasting over UDP
- TCP ping.

If you don't know whether your Central servers can use IP multicasting, you can use TCP ping.

Important: If the Central server has more than one network interface, you must bind the server to a particular network address. You do this in the Central.properties file. For the steps for binding the server to a network address, see the next procedure after this one.

To configure Central cluster nodes for failover and automatic run recovery

1. Start Central.
2. Click the **Administration** tab, and then click that tab's **System Configuration** subtab.
3. Scroll down to the **Clustering Settings** section.

Clustering Settings <input type="checkbox"/> Clustering Enabled	
Description	Value
The class D multicast address to use when UDP is selected as the protocol.	228.10.10.10
The multicast port to use when UDP is selected as the protocol.	45566
The name of the cluster this Central belongs to.	CENTRAL_CLUSTER
The protocol to use to communicate with other Central cluster members (UDP or TCP). UDP specifies to use UDP IP multicasting between cluster nodes. TCP specifies to use TCP communication between cluster nodes.	UDP
A comma separated list of host[port] specifications of all Central hosts and their TCP listener ports in the cluster (i.e. hosta.mycompany.com[45566],hostb.mycompany.com[45567]). This is only used if TCP is selected as the protocol.	localhost[45566]
The TCP listener port to use when TCP is selected as the protocol.	45566

Figure 8 - Creating and configuring a failover and run recovery cluster

4. To enable clustering for failover and automatic run recovery of a failed node's runs by another node in the Central cluster, select the **Clustering Enabled** checkbox.

Notes:

- Automatic run recovery does not require that the Central server is a member of a failover cluster. Clustering enables another node in the cluster to resume the runs immediately, without your having to wait for the failed Central server to recover.
- Only headless runs are recovered if the Central server they were running on crashes while the run was in a Running state.

Next, you'll define the multicast address that the cluster nodes use to communicate with each other.

Each Central cluster requires a unique pair of settings for the cluster's multicast address (mcast_addr) and the multicast port (mcast_port). The values used for the mcast_addr and mcast_port must not overlap with the IP address/port combination used for other applications that run in the environment.

5. In the **Value** box for **The class D multicast address to use when UDP is selected as the protocol** setting, type the IP address to use for multicasting. The address you specify must not to conflict with other applications
6. In the **Value** box for **The multicast port to use when UDP is selected as the protocol** setting, either retain the default (45566) or change it to a port number that:
 - Is not blocked on any of the cluster nodes.
 - Is not used for other purposes in your production environment.

7. In the **Value** box for **The name of the cluster** setting, name the cluster.
8. In the **Value** box for **The protocol to use to communicate with other Central cluster members...** setting, select either UDP or TCP.
9. If TCP is specified for the cluster's internal communication protocol, in the **Value** box for the **A comma-separated list of host[port] specifications of all Central hosts...** setting, type the name-port pairs of the cluster nodes, enclosing each port number in square brackets and separating the pairs with commas.

For instance, if your nodes were:

- edgar.mydomain.ad, using port 888
- rosalind.mydomain.ad, using port 555

Important: The port numbers that you specify must not be blocked on any of the cluster nodes, nor be used for other purposes in your production environment.

You would type the following in the **Value** box:

edgar.mydomain.ad[888],rosalind.mydomain.ad[555]

10. If TCP is specified for the cluster's internal communication protocol, in the **Value** box for the **The TCP listener port to use...** setting, type the port over which the Central node will listen to the other cluster nodes.
11. Click **Save Clustering Settings**.

To bind a Central server to a network address

1. In the Central home directory, open the \conf subdirectory.
2. Locate the Central.properties file and open it in a text editor.
3. Locate the following lines:

```
dharmajgroups.prop.udp=UDP(down_thread=false;mcast_send_buf_size=64000;mcast_port=${clustering.mcast_port:45566};\
discard_incompatible_packets=true;ucast_rcv_buf_size=2000000;mcast_addr=${clustering.mcast_addr:228.10.10.10};\
up_thread=false;loopback=false;mcast_rcv_buf_size=25000000;max_bundle_size=64000;\
max_bundle_timeout=30;use_incoming_packet_handler=true;use_outgoing_packet_handler=false;\
ucast_send_buf_size=640000;tos=16;enable_bundling=true;ip_ttl=2)
```

4. Immediately preceding the end parenthesis, add

```
;bind_addr=<IPADDRESS>
```

where **<IP_ADDRESS>** is the IP address of the network interface that identifies the Central to other machines on the network.

For example, to bind the server to the IP address 255.255.0.225, you would change the lines to read as follows:

```
dharmajgroups.prop.udp=UDP(down_thread=false;mcast_send_buf_size=64000;mcast_port=${clustering.mcast_port:45566};\
discard_incompatible_packets=true;ucast_rcv_buf_size=2000000;mcast_addr=${clustering.mcast_addr:228.10.10.10};\
up_thread=false;loopback=false;mcast_rcv_buf_size=25000000;max_bundle_size=64000;\
max_bundle_timeout=30;use_incoming_packet_handler=true;use_outgoing_packet_handler=false;\
;bind_addr=255.255.0.225)
```

```
ucast_send_buf_size=640000;tos=16;enable_bundling=true;ip_ttl=2;\
bind_addr=255.255.0.225)
```

5. Save your changes and close Central.properties.

Verifying that the Central failover/run-recovery cluster is correctly configured

To verify that the Central failover cluster setup is correct

- Review the Central\logs\wrapper.log for messages such as the following, which indicate the successful formation and configuration of the cluster:

```
INFO | jvm 1 | 2007/08/09 19:26:57 | INFO
[WrapperSimpleAppMain] (19:26:57,309)
com.iconclude.dharma.commons.cluster.JGroupsClusterService - Now
coordinator of cluster group MYCLUSTER_CENTRAL
INFO | jvm 1 | 2007/08/09 19:26:57 | INFO
[WrapperSimpleAppMain] (19:26:57,309)
com.iconclude.dharma.services.cluster.CentralClusterService -
Elected cluster master - starting master services
INFO | jvm 1 | 2007/08/09 19:26:57 | INFO
[WrapperSimpleAppMain] (19:26:57,309)
com.iconclude.dharma.commons.cluster.JGroupsClusterService - Node
10.255.147.110:45566 added to cluster group MYCLUSTER_CENTRAL
INFO | jvm 1 | 2007/08/09 19:28:50 | INFO [UpHandler (TCPPING)]
(19:28:50,373)
com.iconclude.dharma.commons.cluster.JGroupsClusterService - Node
10.255.147.111:45566 added to cluster group MYCLUSTER_CENTRAL
```

OR

To verify the cluster status, use the Get Cluster Servers flow . This flow queries a node in the cluster to retrieve the entire state of the cluster, the master node and all the slave nodes that have been discovered. (To locate the Get Cluster Servers flow, use **Search** in Central.)

Replicating the repository across the nodes of the cluster

Each of the nodes in a cluster has its own repository, and all of their repositories must be identical to each other. Thus, after you have created the cluster, you must replicate the repository from the master node to the slave nodes.

The following procedure assumes that, per best practices, you have a Central server that is not in the production environment (and so is not a member of the production Central cluster), but rather is designated as a staging Central server. You can, however, run the Publish Staging to Production Cluster flow from one of the (target) cluster members or even from a server that is neither the staging server nor a member of the cluster.

Note: The user account that runs the Publish Staging to Production Cluster flow must have sufficient capabilities and permissions on the relevant Central servers (the source server, the target servers, and the server from which the flow is run).

To copy the Central repository across the nodes of the cluster

1. When you are ready to distribute new work to the Central cluster in your production environment, publish the repository to the staging Central server.
2. Run the flow Publish Staging to Production Cluster, providing the URL of the staging server and the URL of one of the Central servers in the cluster.

Tip: To find the Publish Staging to Production Cluster flow, use **Search** in Central.

Here is an example of server URLs that you might supply to the flow inputs:

- If the Staging Central server is named "staging.domain1.ad", its URL would be <https://staging.domain1.ad:8443>
- If Production Central server cluster members are prod1.domain2.ad, prod2.domain2.ad, and prod3.domain2.ad, you could use any of the following URLs for the serverURL input: <https://prod1.domain2.ad:8443>, <https://prod2.domain2.ad:8443>, or <https://prod3.domain2.ad:8443>

Configuring Linux Virtual Server load balancer to use HP OO

To provide load balancing for Central on a Linux Virtual Server

1. Install a Linux Virtual Server and configure it for load balancing.
For more information on setting up and configuring a Linux Virtual Server for load balancing, see <http://www.linuxvirtualserver.org/>.
2. When configuring the load balancer, be sure to configure two ports in the following fashion:

- **With IP affinity for browser connections**

Set the NLB timeout to 31 minutes. This sets the limit before timing out to 31 minutes for any browser, which is a minute longer than the Central session timeout. This prevents a Central user's session from timing out and ensures that the user is directed to same node throughout his or her session (unless of course, the node fails). The affinity is handled through IP. After the timeout, on subsequent requests, the user may or may not be directed to the same node as their last session.

- **Without IP affinity for automatic connections**

This is a useful setting for invoking flows automatically through the use of another tool. To ensure proper load-balancing for a flood of requests from a single machine, requests for automatic connections should not have IP affinity on, and should be requested on a different port than for browser connections.

For example, if IP affinity is turned on for all requests coming from <https://virtualmachine:8443/OpsForce> and IP affinity is turned off for all requests coming from <https://virtualmachine:443/OpsForce>, both requests would be mapped, via the LVS to hit <https://nodeX:8443/OpsForce>.

3. When all your configurations are complete for the Linux Virtual Server, add each Central node behind the HP OO Load-Balancing cluster.
4. To test the cluster by:
 - a. Make requests from different machines via a browser.
The requests should be routed correctly to each node, and automatic requests should also be routed correctly.

- b. Take down each node in succession, then bring them back up to make sure they rejoin the cluster.

Uninstalling HP OO programs

Before uninstalling HP OO programs, be sure to back up your installation and repository. For information on backing up HP Central and Studio, see the HP OO *Administrator's Guide* (AdminGuide.pdf).

When you remove HP OO Central, the RAS is also uninstalled.

Note: It is recommended that, before uninstalling Central and RAS, you manually stop the HP OO Windows services (RSCentral, RSJRAS, and RSScheduler). The uninstall programs try to shut down the HP OO Windows services, but if the system is in heavy use when you uninstall, they may not be able to shut the services down and so may leave files and directories remaining after the uninstallation has finished. If this happens, you can manually delete the files. However, you can also facilitate a clean uninstallation by manually stopping the HP OO services before uninstalling.

To uninstall Studio

1. Make sure Studio is shut down.
2. Open **Control Panel**, then **Add/Remove Programs**.
3. Scroll down to and highlight **HP Operations Orchestration Studio 7.20**, and then click **Remove**.
4. When you are prompted to confirm whether you want to remove Studio and its components, click **Yes**.

The **Uninstall Status** box appears, in which progress of the removal is tracked on a progress bar.

When Studio is completely removed, a message box appears, telling you so.

If you have relevant directories open, the message box might tell you that some components could not be removed. This is not significant.

5. Click **OK**.

To uninstall Central

1. Close any open instances of the Central web application.
2. Open **Control Panel**, then **Add/Remove Programs**.
3. Scroll down to and highlight **HP Operations Orchestration Central 7.20**, and then click **Remove**.
4. When you are prompted to confirm whether you want to remove Central and its components, click **Yes**.

The **Uninstall Status** box appears, in which progress of the removal is tracked on a progress bar.

When Central is completely removed, a message box appears, telling you so.

5. Using **Add/Remove Programs** in the Windows Control Panel, remove **HP Operations Orchestration Central**.

Index

- Central
 - Web application defined, 1
 - Web client defined, 1
- Central
 - installing, 11
- Central
 - uninstalling, 32
- Central cluster
 - configuring, 27
 - creating, 27
- Central server
 - network address, 27
- Cluster
 - configuring, 27
 - creating, 27
- copyright notices, 1
- Database
 - configuring, 9
- Database user
 - configuring, 9
- Default ports, 2
- Failover cluster
 - configuring, 27
 - creating, 27
- Firewalls, crossing, 2
- HP OO
 - crossing firewalls, 2
 - deploying to users, 19
 - installing with clusters, high-level view, 7
 - installing, high-level view, 6
 - overview, 1
 - ports, 2
- HP OO Scheduler
 - uninstalling, 32
- Installation
 - planning, 1
- Installation, basic, 5
- Installation, testing, 19
- Installing, roadmap to, 1
- legal notices, 1
 - copyright, 1
 - restricted rights, 1
 - trademark, 1
 - warranty, 1
- Load balancing
 - on a Linux Virtual Server, 31
- network address
 - binding, 27
- Operations Orchestration. *See* HP OO
- Ports, 2
- RAS
 - standalone installation, 20
 - testing the installation, 22
 - uninstalling, 32
- Remote Action Service. *See* RAS
- restricted rights legend, 1
- Run-recovery cluster
 - configuring, 27
 - creating, 27
- Studio
 - defined, 1
- Studio
 - uninstalling, 32
- Testing installation, 19
- trademark notices, 1
- warranty, 1