



# Opsware™ Process Automation System

*Version 7.0, Installing Opsware Process Automation  
System*

For Process Automation System Administrators

Copyright © 2000-2007 Opware Inc. All Rights Reserved.

Opware Inc. Unpublished Confidential Information. NOT for Redistribution. All Rights Reserved.

Opware is protected by U.S. Patent Nos. 6,658,426, 6,751,702, 6,816,897, 6,763,361 and patents pending.

Opware, SAS Web Client, Model Repository, Data Access Engine, Web Services Data Access Engine, Software Repository, Command Engine, Opware Agent, Model Repository Multimaster Component, and Code Deployment & Rollback are trademarks and service marks of Opware Inc. All other marks mentioned in this document are the property of their respective owners.

Additional proprietary information about third party and open source materials can be found at <http://www.opware.com/support/sas65tpos.pdf>.

# Table of Contents

<b>Table of Contents</b> .....	<b>3</b>
<b>Updating documentation</b> .....	<b>4</b>
<b>Overview</b> .....	<b>4</b>
<b>Roadmap to this guide</b> .....	<b>5</b>
<b>Planning the installation</b> .....	<b>5</b>
Useful information .....	6
Ports of importance to PAS components .....	6
Crossing firewalls .....	6
Server and Studio distribution .....	6
One staging Central server, one production Central server .....	7
Multiple authors .....	7
Multiple authors, multiple Central production servers .....	8
Extensibility .....	8
Scheduling automatic runs of flows .....	8
Creating a security certificate .....	8
<b>Basic PAS installation</b> .....	<b>9</b>
High-level view of installing PAS .....	9
High-level view of installing PAS with clustering .....	10
Configuring the database and database user .....	13
Schema and schema user configuration with Oracle 10g Enterprise .....	13
Database and database user configuration with MySQL .....	13
Database and database user configuration with SQL Server .....	13
Installing Central .....	14
Installing Studio .....	17
Installing Scheduler .....	19
<b>Testing the installation</b> .....	<b>21</b>
<b>Making PAS available to users</b> .....	<b>21</b>
<b>Installing JRAS and/or NRAS on standalone servers</b> .....	<b>22</b>
Installing JRAS .....	22
Installing the .NET Remote Action Service .....	23
<b>Clusters and PAS</b> .....	<b>24</b>
Creating a load-balancing cluster for PAS Central .....	26
Installing the PAS Load Balancer .....	27
Creating and configuring a Central cluster for load-balancing .....	28
Creating a failover cluster for PAS Load Balancer .....	29

Creating a Central cluster for failover and run recovery.....	29
Replicating the repository across the nodes of the cluster.....	33
Configuring Linux Virtual Server load balancer to use PAS.....	33
<b>Uninstalling PAS programs.....</b>	<b>34</b>
<b>Index .....</b>	<b>35</b>

## Updating documentation

Documentation enhancements are a continual project at Opsware. You can update the documentation set at any time using the following procedure (which is also available in the PAS readme file).

### To obtain PAS documentation

1. On The Opsware Network Web site (<https://support1.opsware.com/support/index.php>), log in with the Opsware Network account name and password that you received when you purchased PAS.
2. On the **Support** tab, click the **Product Docs** subtab.
3. Under **Quick Jump**, click **Process Automation System**.
4. Under **Process Automation System**, click **ZIP** beside **PAS 7.0 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**, and then 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 Opsware\PAS home directory, in either the **\Central\docs\help\Central** or **\Studio\docs\help\Studio** subdirectory, overwriting the existing file.

## Overview

Opware Process Automation System (PAS) is a system for creating automated procedures (called *Ops 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.

Process Automation System (PAS) is made up of the PAS Central Web application, the Central Web client, and PAS Studio:

- The PAS Central Web application provides the business-logic tier between the PAS Repository database and both the PAS Central Web client and PAS Studio.

- The PAS Central Web client is the browser-based tool that IT personnel use to execute Ops flows.
- PAS Studio is a standalone application in which you create new Ops flows.

## Roadmap to this guide

Installing PAS 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 PAS extensibility.

**Important:** To run the Windows Accelerator Pack Ops flows against a remote machine, you must install Java Remote Action Service (JRAS) and .NET Remote Action Service (NRAS) on the remote machine.

- [Basic PAS installation](#)

Basic installation includes the Central Web application and Studio.

- [Making PAS 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 PAS](#)

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

For information on ongoing administrative tasks and on configuring PAS for extended functionality, see the *PAS Administration Guide*.

## Planning the installation

Planning issues include the following:

- Server and Studio distribution
- Ports for crossing firewalls
- Extensibility
- Scheduling automatic runs of flows
- Clustering for scalability and high availability:
  - Whether to cluster PAS Central servers for load-balancing and/or failover support
 

If you create a load-balancing cluster for Central servers, whether you will use PAS 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 PAS Central database servers
  - Whether to cluster computers on which JRAS and NRAS are installed
  - Whether the Central server nodes support IP multicasting
 

If not, then the Central nodes use TCP Ping.

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

## Useful information

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

## Ports of importance to PAS components

By default, PAS 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 PAS components, such as Central, Scheduler, JRAS, and NRAS: 18443
- JRAS: 9004
- NRAS: 9005
- Scheduler: 19443
- SMTP server: 25

## Crossing firewalls

When you install NRAS or JRAS 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 NRAS or JRAS. For standalone JRAS, this is 9004; for standalone NRAS, this is 9005.

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 NRAS and/or JRAS on a standalone server. On each remote computer where you need standalone installations of NRAS or JRAS, you can install one of each. You cannot install a standalone NRAS or JRAS on the Central Web server; the Central installation program includes installation of NRAS and JRAS.

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:

## 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.

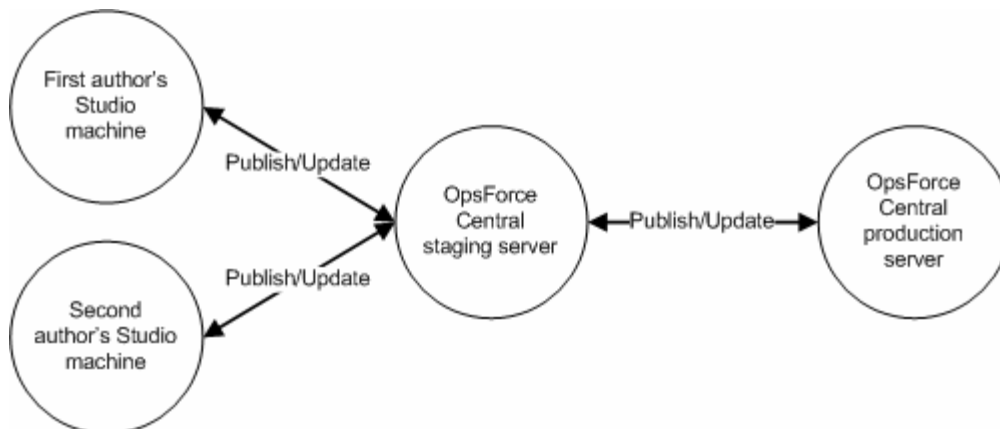


**Figure 1 - Minimum recommended configuration**

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

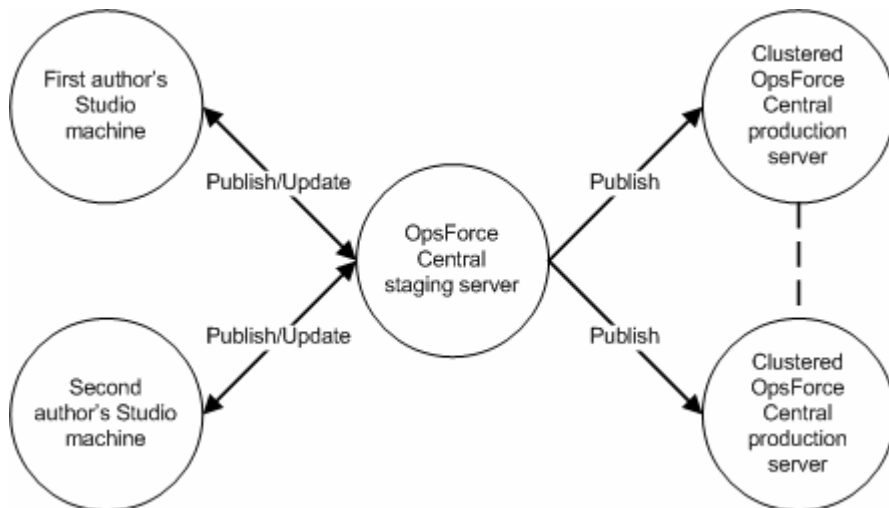
Multiple authors can exchange work by publishing and updating from the Central staging server. When either author publishes to the staging server, his or her work is copied there. When either author updates from the staging server, his or her local repository is updated with any work that the other author 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 PAS operations, you can carry out Ops flows that are enabled through the Web service on which runs. Using JRAS or NRAS, you can extend the ability of PAS 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.

## Scheduling automatic runs of flows

Doing so requires installing the PAS Scheduler.

## 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 PAS 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 PAS installation

Installing Central and Studio provides everything you need to create and run Ops flows. The Central installation includes automatic installation of Java Remote Action Service (JRAS) and .NET Remote Action Service (NRAS), which enable you to run Ops flows that operate remotely from PAS or interact with other programs.

## High-level view of installing PAS

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 JRAS or NRAS installation.

You would install and configure PAS components in the following order:

1. Make sure that the machines on which you install the PAS components meet hardware and software requirements described in:
  - The Opsware PAS System Requirements document (PAS\_SystemReqs.pdf)
  - Requirements for all the required software programs.
2. Create a testing/staging PAS environment, which we will assume has the following configuration:
  - Database server
  - Central server
  - Optional: installation of JRAS and/or NRAS each on its own computer

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

- d. Install the database management system on the database server.
- e. Create a database for Central and configure the database and database user as necessary.

For information, see [Configuring the database and database user](#).

- f. Install Central.

For information, see [Installing Central](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Central folder.

- g. Install Studio.

For information, see [Installing Studio](#).

- h. Install Central Scheduler.

For information, see [Installing Scheduler](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Scheduler folder.

- i. If necessary, install standalone JRAS and/or NRAS.  
For information, see [Installing JRAS and/or NRAS on standalone servers](#).

OR

If you are installing a standalone JRAS on a Linux system, see README.linux.txt in the Linux/JRAS folder.

3. Create the production PAS 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 JRAS and/or NRAS 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, see [Installing Central](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Central folder.

- e. Install Central Scheduler.  
For information, see [Installing Scheduler](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Scheduler folder.

- f. Install Studio on each author's machine, referencing the Central installation.  
For information, see [Installing Studio](#).
- g. If necessary, install standalone JRAS and/or NRAS (clustered, thus more than one).  
For information, see [Installing JRAS and/or NRAS on standalone servers](#).

OR

If you are installing a standalone JRAS on a Linux system, see README.linux.txt in the Linux/JRAS folder.

## High-level view of installing PAS 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 JRAS or NRAS installation.

You would install and configure PAS components in the following order:

1. Make sure that the machines on which you install the PAS components meet hardware and software requirements described in:
  - The Opware PAS System Requirements document (PAS\_SystemReqs.pdf)
  - Requirements for all the required software programs.
2. Create the following testing/staging PAS environment, which we will assume does not need load-balancing or failover support:

- Database server
- Central server
- Optional: installation of JRAS and/or NRAS each 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, see [Installing Central](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Central folder.

- d. Install Studio.

For information, see [Installing Studio](#).

- e. Install Central Scheduler.

For information, see [Installing Scheduler](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Scheduler folder.

- f. If necessary, install standalone JRAS and/or NRAS.

For information, see [Installing JRAS and/or NRAS on standalone servers](#).

OR

If you are installing a standalone JRAS on a Linux system, see README.linux.txt in the Linux/JRAS folder.

3. Create the following production PAS 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 JRAS and/or NRAS 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, see [Installing Central](#).

OR

If you are installing on a Linux system, see README.linux.txt in the Linux/Central folder.

- e. Install Central Scheduler.  
For information, see [Installing Scheduler](#).
- OR
- If you are installing on a Linux system, see README.linux.txt in the Linux/Scheduler folder.
- f. Install PAS Load Balancer (probably on one of the Central servers).  
For information, see [Installing the PAS Load Balancer](#), which includes alternative instructions for installing on either a Linux or on a Windows system.
- g. Using the Load Balancer, create and configure the Central load-balancing cluster.  
For information, see [Creating and configuring a Central cluster for load-balancing](#).
- h. 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](#).
- i. Create a failover cluster for PAS Load Balancer.  
For information, see [Creating a failover cluster for PAS Load Balancer](#).
- j. Install Studio on each author's machine, referencing the Central installation.  
For information, see [Installing Studio](#).
- k. If necessary, install standalone JRAS and/or NRAS (clustered, thus more than one).  
For information, see [Installing JRAS and/or NRAS on standalone servers](#).

OR

If you are installing a standalone JRAS on a Linux system, see README.linux.txt in the Linux/JRAS folder.

- l. Using the Load Balancer, create and configure load-balancing clusters for the JRAS and/or the NRAS installations.  
The method is the same as for creating a Central load-balancing cluster.
- m. 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

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

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

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.

### Schema and schema user configuration with Oracle 10g Enterprise

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

### Database and database user configuration with MySQL

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

### Database and database user configuration with SQL Server

If another person installs PAS 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 PAS 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.
  - n. Within the SQL Server Configuration Manager program, expand **SQL Server 2005 Network Configuration**.
  - o. 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 PAS 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, if you wish, the Scheduler) and Studio.

## Installing Central

The PAS Central installation checks connectivity with and configures the database that you created and installs the Central server files.

**Note:** By default, Central is configured to communicate with remote installations of JRAS and NRAS over ports 9004 and 9005, respectively (which the JRAS and NRAS installations configure 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, you will need to obtain the MySQL driver file.

### To install PAS Central

1. Before starting the installation, make sure that the SQL Server is connected to the database that you're using for .
2. If Central will use a MySQL database, obtain the MySQL driver file Mysql-connector-java-<versionnumber>.jar and copy it to a location of your choosing.
3. Navigate to and double-click PASCentralInstaller-sqlserver.exe.
4. On the welcome page of the PAS Central Setup Wizard, click **Next**.
5. On the **License Agreement** page, read the agreement, select **I accept the agreement**, and then click **Next**.

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

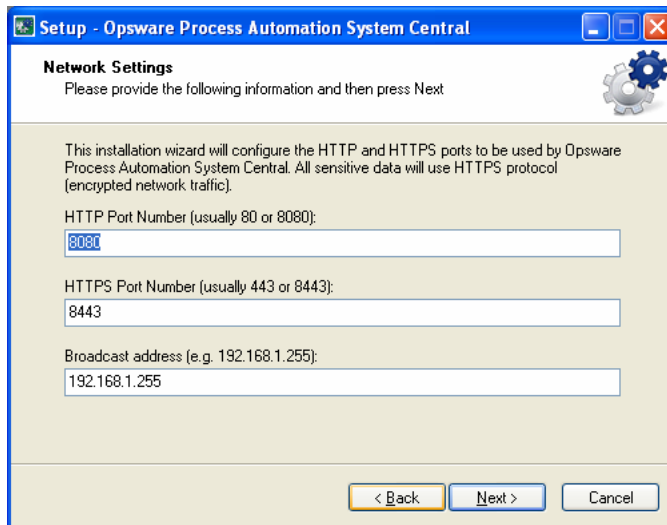
OR

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

If you are reinstalling the Web application or have already installed Studio to the same directory, a message warns you that the directory already exists and prompts you to confirm that you want to overwrite files in the directory.

7. In the message, click **Yes**.

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



The screenshot shows a window titled "Setup - Opware Process Automation System Central". The main heading is "Network Settings" with a gear icon. Below the heading, it says "Please provide the following information and then press Next". A paragraph explains: "This installation wizard will configure the HTTP and HTTPS ports to be used by Opware Process Automation System Central. All sensitive data will use HTTPS protocol (encrypted network traffic)." There are three input fields: "HTTP Port Number (usually 80 or 8080):" with the value "8080", "HTTPS Port Number (usually 443 or 8443):" with the value "8443", and "Broadcast address (e.g. 192.168.1.255):" with the value "192.168.1.255". At the bottom, there are three buttons: "< Back", "Next >", and "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.

8. Either accept the default port numbers or type different port numbers in the text boxes
9. 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.
10. Either accept the default broadcast IP address or type a different one, and then click **Next**.

The **Database Information** page appears.

**Setup - Opsware Process Automation System Central**

**Database Information**  
 Opsware Process Automation System Central needs configuration information about the Database. Please ask the system administrator if you do not know these

DB Host or IP Address: localhost

DB Port Number: 1433

DB Name:

DB User Name:

DB Password:

Confirm DB Password:

Create (or recreate, if it exists) the Database?

**DB Name format:** <db\_name>[instance=<instance\_name>]  
**DB User Name format:** [<domain\_name>]<db\_user\_name>

< Back   Next >   Cancel

**Figure 5 - Database Information page**

If you are upgrading from version 2.2, do not recreate the database. There are no changes to the database schema between versions 2.2 and 7.0, and if you recreate the database, you lose historical data.

**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>
- 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.

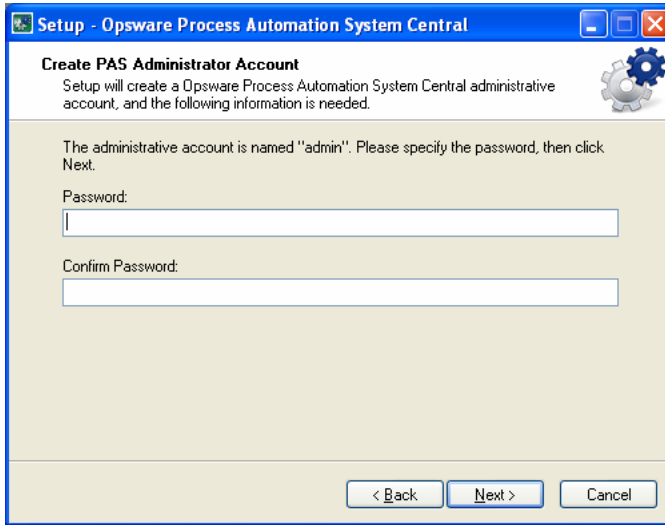
- Password and Confirm Password  
 The database name (**DB Name**) and password are the name and password that the Central Web application uses to access the database.

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

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



The **Create PAS Administrator Account** page appears. The administrator account, which you will use to log in to Central, is always named "admin." In this dialog box, you create a password for the admin account.



**Figure 6 - Create PAS Administrator Account page**

13. In the **Password** and **Confirm Password** boxes, type a password.  
The **Ready to Install** page appears, displaying the information that you specified in the rest of the setup.
14. To proceed, click **Install**.  
If you are installing using a MySQL database, the **Open File** page appears.
15. If you are installing using a MySQL database, browse to the MySQL driver file (Mysql-connector-java-<versionnumber>.jar, and then click **Next**.  
The installation program tracks progress on the **Installing** page.
16. When the installation completes, click **Finish**.  
The remaining programs to install are Studio and Scheduler. Installing Scheduler installs a **Schedule** tab in Central. The installation procedure for Scheduler follows the section on 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 PAS Studio Setup Wizard, navigate to and double-click PASStudioInstaller.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, type the location where you want the Studio files installed.

OR

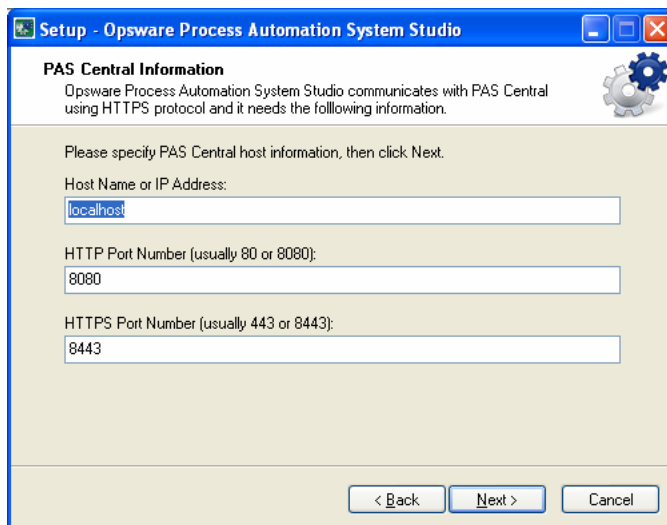
Click **Browse**, specify a different location where you want the Web application installed, and then click **Next**.

If you install the Studio on the same machine as the Central Web application, make sure that the location that you specify is the same location in which the Web application is installed.

If you have already installed Studio or the Central Web application, a message appears, warning you that the directory already exists and prompts you to confirm that it is acceptable to overwrite files in the directory.

5. In the message, click **Yes**.

The **PAS Central Information** page appears.



**Figure 7 - PAS Central Information page**

6. 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.

7. 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.

8. 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.

9. To proceed, click **Install**.

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

10. When the installation completes, make sure that the **Launch PAS Studio** box is selected, and then click **Finish**.

Starting Studio is part of testing the success of the installation.

11. 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.

## Installing Scheduler

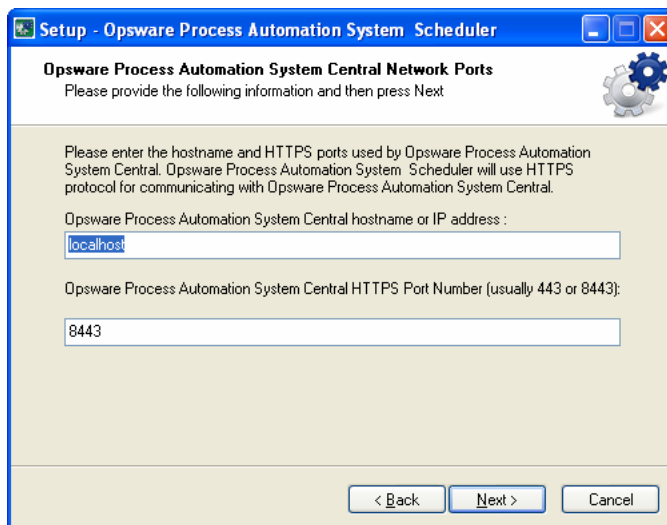
**Important:** If you are installing using a MySQL database, you will need to obtain the MySQL driver file.

### To install PAS Scheduler

1. If Scheduler will use a MySQL database, obtain the MySQL driver file `Mysql-connector-java-<versionnumber>.jar` and copy it to a location of your choosing. You probably already completed this step when installing Central.
2. Navigate to and double-click `PASSchedulerInstaller-sqlserver.exe`.
3. On the **Welcome** page, click **Next**.
4. On the **License Agreement** page, accept the terms of the license agreement, 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 Studio files installed.
  - Click **Browse** and specify a different location.

If you have already installed Studio or the Central Web application, a message appears, warning you that the directory already exists and prompts you to confirm that it is acceptable to overwrite files in the directory.

The **PAS Central Network Ports** page appears.



**Figure 8 - PAS Central Network Ports dialog box**

6. Specify the following, and then click **Next**:
  - The host server on which the Central Web application is installed
  - The HTTPS port number specified when you installed Central
7. In the **Enter PAS Scheduling Account** dialog, type the name of a valid PAS account that has sufficient capabilities and access rights to schedule and start flows automatically.

It is simplest to specify the admin account here. The admin account possesses every capability and all access rights to all flows. (Capabilities are the objects that allow a group to perform specific functions; authors assign groups access rights to specific flows. For more on Capabilities and access rights, see the Help systems for Central and for Studio.)

If you choose to specify another account, keep in mind that if the account is not a valid PAS account (either a valid AD account or an PAS file-based account), an “HTTP 401 Unauthorized” error will occur when the scheduler tries to run a scheduled flow.

The **Database Information** dialog box appears.

**Figure 9 - Scheduler Database Information dialog box**

8. Fill in the information as your configuration dictates, specifying the same database and database user credentials that you did when installing Central. 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>`
  - 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
9. To create a new database, select the **Create (or recreate, if it exists) database** 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.
  10. To proceed, click **Install**.  
If you are installing using a MySQL database, the **Open File** page appears.
  11. If you are installing using a MySQL database, browse to the MySQL driver file (Mysql-connector-java-<versionnumber>.jar, and then click **Next**.  
The Setup Wizard tracks progress on the **Installing** page.
  12. When the installation completes, click **Finish**.

## Testing the installation

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 PAS available to users

For PAS 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, PAS, make sure that client machines meet hardware and software requirements for Central or Studio, as described in PAS\_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/static/Login.htm>

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 PAS ADMINISTRATOR role.

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

- If you have not installed the author's copy of Studio, PASStudioInstaller.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 PAS ADMINISTRATOR role.

## Installing JRAS and/or NRAS on standalone servers

You can install one installation each of .NET remote action service (NRAS) and Java remote action service (JRAS) on a computer that is remote from the Central server. You cannot install a standalone NRAS or JRAS on the Central Web server, because the Central installation program includes installation of NRAS and JRAS.

**Note:** By default, the standalone JRAS installation configures it to communicate over port 9004, which is also configured on the Central server for communicating with JRAS. In the same way, the standalone NRAS installation configures it to communicate over port 9005, which is also opened on the Central server.

## Installing JRAS

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

### To install the JRAS server

1. On the Opsware PAS CD, locate JRASSetup.exe and copy it to the machine on which you're going to install it.

2. Start the copy of JRASSetup.exe that you just created.

The Opsware JRAS Setup Wizard starts with the **Welcome** page.

3. Click **Next**.

4. On the **License Agreement** page, read the agreement, click **I accept the agreement**, and then click **Next**.

The **JRAS Configuration** page appears. On this page, you specify:

- A location for the Java Virtual Machine that this wizard creates.
- A communication port number for JRAS.

5. In the **JVM Path** box, leave the default installation path.

OR

Specify a different location where Java.exe resides.

6. In the **Port Number** box, leave the default port number or specify a new one, and then click **Next**.

The **JRAS Optional Packages** page appears.

7. Select either or both of the following items of optional content for installation:
  - Base Content, which comprises the IActions that are used in the Opsware Accelerator Packs.
  - SDK Examples, which contain sample code you use to build your own IActions.
8. Click **Next**.
9. On the **Select Destination Location** page, click **Next** to accept the default installation path for JRAS.  
OR  
Click **Browse**, specify a different location where you want JRAS installed, and then click **Next**.
10. On the **Select Start Menu Folder** page, either accept the default **Start** menu folder in which you want the JRAS shortcut to reside or specify a different one, and then click **Next**.  
The **Ready to Install** page appears, displaying the location where JRAS will be installed and other information on the choices you have made.
11. To proceed, click **Install**.  
The installation program tracks progress on the installation progress page.
12. When the installation completes, click **Finish**.

#### To test the JRAS server installation

- Using a Web browser, access the following URL to confirm that the Web service is running.  
<http://<Hostname>:<Portnumber>/JRAS/services/RCAgentService?wsdl>  
where
  - <Hostname> is the name of the PAS Web application server
  - <Portnumber> is the name of the default port used for JRAS on the Web application server

If you accept the default port number in the JRAS Setup, this number is 4085.

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

## Installing the .NET Remote Action Service

The .NET Remote Action Service (NRAS) installation program creates and places files in the default IIS Web root directory.

Because NRAS is installed as part of the Central installation, you only need to perform the NRAS installation on machines on which you have not installed Central.

#### To install NRAS

1. At the source of the Opsware installation files, navigate to and run PASNRASInstaller.exe.
2. On the **Welcome** page, click **Next**.
3. On the **License Agreement** page, read the agreement, click **I accept the agreement**, and then click **Next**.
4. On the **NRAS Optional Packages** page, select either or both of the following items of optional content for installation:

- **Base Content**, which comprises the IActions that are used in the Opaware Accelerator Packs.
  - **SDK Example**, which contain sample code you use to build your own IActions.
5. On the **Network Settings** page, specify an IP address that NRAS will use for communications.
  6. On the **Select Destination Location** page, click **Next** to accept the default installation path for NRAS.

OR

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

7. On the **Select Start Menu Folder** page, either accept the default **Start** menu folder in which you want the NRAS shortcut to reside or specify a different one, and then click **Next**.

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

8. To proceed, click **Install**.  
The installation program tracks progress on the installation progress page.
9. When the installation completes, click **Finish**.

#### To test the NRAS server installation

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

<http://<Hostname>:<Portnumber>/NRAS/services/RCAgentService.asmx>

where

- <Hostname> is the name of the PAS Web application server
- <Portnumber> is the name of the default port used on the Web application server.

If you accepted the default port number when you installed the NRAS server, the default port number is 4080.

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

#### To obtain access to the NRAS interface within Studio

- Open Studio and configure a remote action service.  
For information on configuring a remote action service, see Help for Studio.

## Clusters and PAS

If you run large volumes of PAS flows and if high availability is a requirement for your PAS installation, you can create a cluster that provides your PAS configuration with load-balancing and failover clustering support in any combination of the following:

- A load-balancing cluster for PAS Central servers and any standalone JRAS and NRAS installations

You can use PAS Load Balancer to create, configure, and run such a load-balancing cluster.



- Failover support for the PAS 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 PAS 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 (if you install Scheduler), and system configuration, such as enabling of user authentication providers.  
You can use any clustering software to cluster the Central database servers.

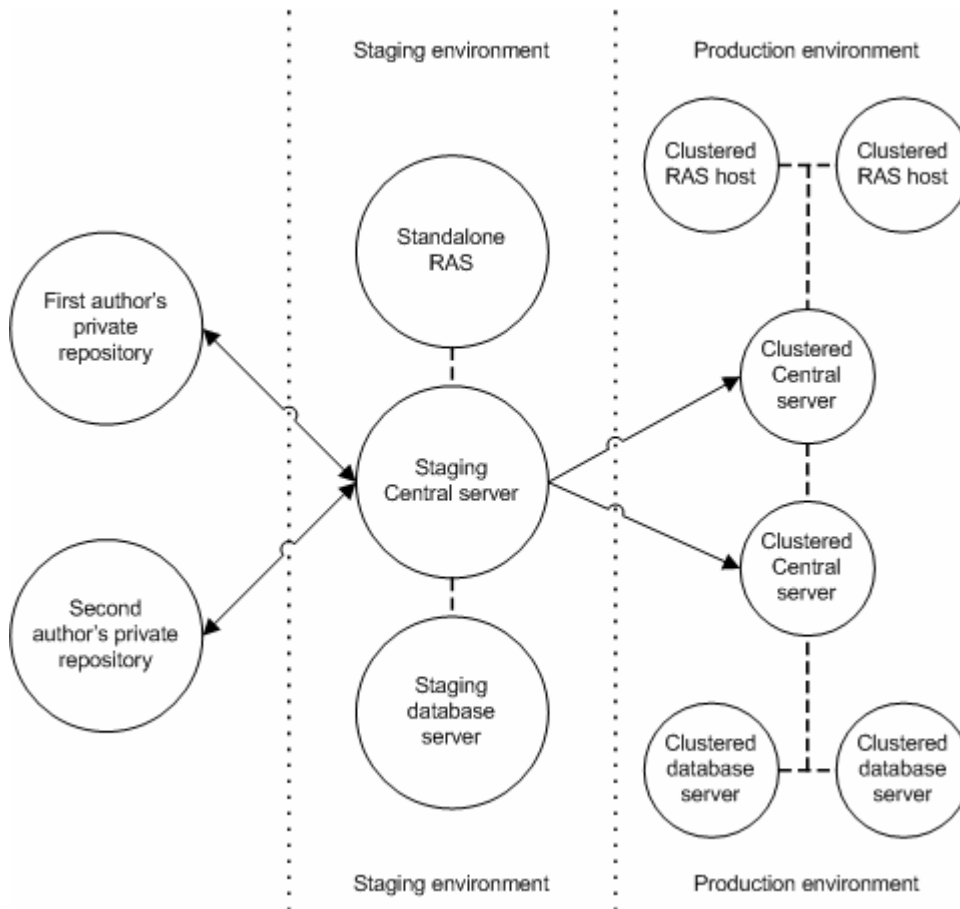
Each PAS Central node must be its own dedicated machine, and use its own repository (a hierarchical folder structure of XML files). PAS provides flows or operations with which you can synchronize repositories and iActions between all the nodes of the application cluster.

To combine multiple-authoring with a PAS Central cluster, install PAS Central in a staging environment to which the authors will publish and from which they will update. This staging instance of PAS Central should be the only source for publishing flows, iActions, and other PAS 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 PAS 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 PAS 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 (JRAS or NRAS) is installed on a standalone host.



**Figure 10 - PAS configuration with staging environment and clustered production environment**

## Creating a load-balancing cluster for PAS Central

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

1. Installing the PAS Load Balancer
2. Creating and configuring a Central cluster for load-balancing
3. Creating a failover cluster for the PAS 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 PAS programs might look like the following:

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

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

Server 1 Load Balancer would handle all HTTP requests and route them to Central 1 or Central 2, JRAS 1 or JRAS 2, NRAS 1 or NRAS 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, JRAS 2, and NRAS 2, as long as Server 1 is down.

## Installing the PAS Load Balancer

The PAS 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 PAS Load Balancer for PAS Central on a Windows operating system

1. To start the PAS Studio Setup Wizard, navigate to and double-click PASClusterInstaller.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 PAS Cluster to be installed.
  - Click **Browse**, specify a different location.
5. In the **Select Start Menu Folder** page, do one of the following, and then click **Next**:
  - Accept the default folder.
  - Type the folder where you want the Studio files installed.
  - Click **Browse** and 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**.

### To install the PAS Load Balancer for PAS Central on a Linux operating system

1. On the installation CD, in the Clustering subdirectory, locate and copy the compressed tar file PASCluster-7.0.tar.gz and uncompress it (`tar -xzf`

PASCluster-7.0.tar.gz) to the location where you're going to perform the installation.

**Note:** The following instructions reproduce the information in INSTALL.txt (contained in PASCluster-7.0.tar.gz).

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 PAS Load-Balancer.

### To configure a cluster

1. Open a command window.
2. Navigate to the location where you installed the PAS 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 (PAS Central, JRAS, or NRAS), 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:      PAS
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 PAS Load Balancer

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

### To create a failover cluster for PAS Load Balancer

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

In the PAS Central Central.properties file, 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.

Enabling automatic run-recovery involves some manual configuration of the Central.properties file on each member of the Central cluster, including:

- Enabling clustering support.

- Enabling run recovery.
- Giving the failover cluster a unique name.

**Important:** JGroups is the mechanism that the Central cluster manages its nodes and with which the nodes recognize each other as members of the cluster. By default, Central uses IP multicasting over UDP for communication between nodes in a cluster. In the step in which you configure the cluster using the JGroups framework, you configure the cluster nodes to communicate with each other using either IP multicasting or TCP ping. If you don't know whether your Central servers can use IP multicasting, you can use TCP ping.

### To configure Central cluster nodes for failover and automatic run recovery

1. Locate the Central.properties file in the PAS\Central home directory and open it in a text editor.
2. In Central.properties, locate the line `dharmia.cluster.support=` and set it to true, as follows:

```
dharmia.cluster.support=true
```

3. To enable automatic run recovery support, locate the line `dharmia.cluster.run_recovery=` and set it to true, as follows:

```
dharmia.cluster.run_recovery=true
```

This line allows runs to be recovered by other Central nodes in the same cluster.

#### Notes:

- Automatic run recovery does not require that the Central server is a member of a failover cluster. That is, you can set this line to true even if `dharmia.cluster.support=false`.
- 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 use JGroups configuration to define the multicast address that the cluster nodes use to communicate with each other.

You configure JGroups with the `dharmia.jgroups.prop` settings in Central.properties. 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.

In Central.properties, there is a commented-out example of `dharmia.jgroups.prop` which shows how to use TCP pinging. The `TCP(start_port=45566)` identifies which port is locally used and should be connected by from other Central nodes in the cluster. This configuration requires that each Central server in the cluster know the hostname and port used by all the other Central nodes in the cluster. `TCPPING(initial_hosts=.....)` is a comma delimited list of `host[port]` values which identify all the other Central nodes in the cluster and what their `start_port` is set to.

If the Central cluster nodes cannot use IP multicasting, the JGroups framework can use TCP pinging for communication between nodes. The advantage of IP multicasting is that it requires less maintenance than TCP ping: When you add a node to a failover Central cluster that uses TCP ping, you must add the node's entry (see `TCPPING(initial_hosts=` in the following step) to every other node's Central.properties file.

4. To configure the JGroups framework for UDP multicasting

- a. In Central.properties, in the line that starts `dharmajgroups.prop=`, locate `mcast_port=45566`.
- b. Either retain the default (45566) or change it to a port number that 1) is not blocked on any of the cluster nodes and 2) is not used for other purposes in your production environment.
- c. Locate `mcast_addr=228.10.10.10` and either retain the default (228.10.10.10) or change it to an address that the cluster should use in your production environment in order not to conflict with other applications.

OR to configure the JGroups framework for TCP ping:

- a. Comment out the line that starts `dharmajgroups.prop=`
- b. Locate the commented-out group of lines that starts `dharmajgroups.prop=TCP` and remove the comment-out characters.

The lines should look as follows:

```
dharmajgroups.prop=TCP(start_port=45566):\
```

```
TCPPING(initial_hosts=host1.mycompany.com[45566],host2.mycompany.com[45566];port_range=5;timeout=3000;num_initial_members=3;up_thread=true;down_thread=true):\
```

```
MERGE2(max_interval=10000;down_thread=false;min_interval=5000;up_thread=false):\
```

```
FD(timeout=2000;max_tries=3;down_thread=false;up_thread=false):\
```

```
VERIFY_SUSPECT(timeout=1500;down_thread=false;up_thread=false):\
```

```
pbcast.NAKACK(max_xmit_size=60000;down_thread=false;use_mcast_xmit=false;gc_lag=0;\
```

```
discard_delivered_msgs=true;up_thread=false;retransmit_timeout=100,200,300,600,1200,2400,4800):\
```

```
pbcast.STABLE(stability_delay=1000;desired_avg_gossip=50000;max_bytes=400000;down_thread=false;up_thread=false):\
```

```
VIEW_SYNC(down_thread=false;avg_send_interval=60000;up_thread=false):\
```

```
pbcast.GMS(print_local_addr=true;join_timeout=3000;down_thread=false;join_retry_timeout=2000;up_thread=false;shun=true):\
```

```
FC(max_credits=2000000;down_thread=false;up_thread=false;min_threshold=0.10):\
```

```
FRAG2(frag_size=60000;down_thread=false;up_thread=false):\
```

```
pbcast.STATE_TRANSFER(down_thread=false;up_thread=false)
```

- c. Within the lines, locate `TCPPING(initial_hosts=host1.mycompany.com[45566],host2.mycompany.com[45566]...` and substitute your cluster nodes' names or IP addresses and, within the square brackets, the port that they will use, separating the node entries with a comma and with the semicolon following the last entry.

For instance, suppose your nodes are:

- edgar.mydomain.ad (IP address of 224.0.10.159), using port 888
- richardii.mydomain.ad (IP address of 224.0.10.260), using port 777
- rosalind.mydomain.ad (IP address of 224.0.10.361), using port 555
- fool.mydomain.ad (IP address of 224.0.10.462), using port 444

then if you added the nodes using their domain names, the line would look like the following:

```
TCPPING(initial_hosts=edgar.mydomain.ad[888],richardii.mydomain.ad[777],rosalind.mydomain.ad[555],fool.mydomain.ad[444];
```

If you added the nodes using their IP addresses, the line would look like the following:

```
TCPPING(initial_hosts=224.0.10.159[888],224.0.10.260[777],224.0.10.361[555],224.0.10.462[444];
```

**Important:** Make sure that the port number that you specify in these entries is 1) not blocked on any of the cluster nodes and 2) not used for other purposes in your production environment.

The node addresses can be any Class D address (in the range 224.0.0.0 to 239.255.255.255) that is unique within your network environment.

5. To give the Central cluster a unique name, in Central.properties, find the following line:

```
dharmacentral_cluster_name=
```

and add =<your\_Central\_cluster\_name>

If, for example, you name the cluster **STAGING**, you would change the line to read,

```
dharmacentral_cluster_name=STAGING
```

6. Save and close Central.properties.

### 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 PAS 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 PAS 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 PAS

### To provide load balancing for PAS Central on a Linux Virtual Server

3. 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/>.
4. 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 PAS 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>.

5. When all your configurations are complete for the Linux Virtual Server, add each PAS Central node behind the PAS Load-Balancing cluster.
6. 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 PAS programs

You uninstall Central, Studio, and JRAS and NRAS by running uninstall programs.

**Note:** It is recommended that, before uninstalling Central and associated programs, you manually stop the PAS Windows services (RSCentral, RSJRAS, RSNRAS, and RSScheduler). The uninstall programs try to shut down the PAS 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 uninstallation has finished. If this happens, you can manually delete the files. However, you can also facilitate a clean uninstallation by manually stopping the PAS services before uninstalling.

### To uninstall JRAS

7. From the **Start** menu, point to **All Programs, Opsware, and PAS**, and then click **Uninstall Opsware JRAS**.
8. When you are prompted to confirm that you want to remove JRAS and its components, click **Yes**.

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

When JRAS 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.

9. Click **OK**.

## To uninstall NRAS

1. From the **Start** menu, point to **All Programs, Opsware, and PAS**, and then click **Uninstall Opsware NRAS**.

2. When you are prompted to confirm that you want to remove NRAS and its components, click **Yes**.

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

When NRAS 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.

3. Click **OK**.

## To uninstall PAS Studio

1. From the **Start** menu, point to **All Programs, Opsware, and PAS**, and then click **Uninstall PAS Studio**.

2. 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.

3. Click **OK**.

## To uninstall PAS Scheduler

- Using **Add/Remove Programs** in the Windows Control Panel, remove **Opsware Process Automation System Scheduler**.

## To uninstall PAS Central

- Using **Add/Remove Programs** in the Windows Control Panel, remove **Opsware Process Automation System Central**.

# Index

## Database

configuring, 14

## Database user

configuring, 14

Default ports, 6

Firewalls, crossing, 6

## Installation

planning, 6

Installation, basic, 9

Installation, testing, 22

Installing, roadmap to, 5

## JRAS

installing, 23

standalone installation, 23

testing the installation, 24, 25

uninstalling, 35

## Load balancing

on a Linux Virtual Server, 34

NRAS

- installing, 24
- standalone installation, 23
- uninstalling, 35

PAS

- crossing firewalls, 6
- deploying to users, 22
- installing with clusters, high-level view, 11
- installing, high-level view, 10
- overview, 5
- ports, 6

PAS Central

- Web application defined, 5
- Web client defined, 5

PAS Central

- installing, 15

PAS Central

- uninstalling, 35

PAS Scheduler

- installing, 20
- uninstalling, 35

PAS Studio

- defined, 5

PAS Studio

- uninstalling, 35

Ports, 6

Process Automation System. *See* PAS

Scheduler. *See* PAS Scheduler

Testing installation, 22