Administrator's Guide

Netscape Application Server:

Process Automation Edition

Version 4.0

Recycled and Recyclable Paper

Version 4.0

Part Number 151-07664-00

Printed in the United States of America.

Netscape Communications Corporation, 501 East Middlefield Road, Mountain View, CA 94043

Contents

Introduction 7
About This Guide 7
Conventions Used in This Guide 7
Viewing Documentation Online 8

Chapter 1 Introduction to PAE Administration 9
Overview 10
Directories in PAE 12
Netscape Application Servers 17
PAE Applications 17
Security in PAE 18
Process Administrator 18
Cluster Management Forms 19
Application Forms 19
Directory Server Terms and Attributes 19
LDAP Terms 20
LDAP Attributes 21
Directory Structure 22
LDAP URLs 22

Chapter 2 Configuration Files 25
Overview 26
The Enterprise Server Files 26
The Process Builder Preferences File 27

PAE Files and Folders 28
Error Messages and Logs 28

Chapter 3 Clusters 31
About Clusters 32
Creating a Cluster 33
New Cluster Information 34
What Happens at Cluster Creation 39
Joining an Existing Cluster 40
Existing Cluster Information 40
What Happens at Cluster Joining 42
Managing a Cluster 42
Changing Cluster Information 43
Accessing Directory Server Information 45
Viewing the Process Administrator Logs 46
Unjoining From a Cluster 47
Deleting a Cluster 47

Chapter 4 Applications 51
About Applications 51
Stopping & Starting Applications 54
Autostart 55
Initializing an Application 55
Stop 56
Start 56
Closing Applications 57
Uninstalling Applications 58
Viewing the Application Logs 60
Exporting and Deleting Data 61

Chapter 5 Databases 65
Overview 66
Using Databases With PAE 67
Cross-Application Tables 68
Primary & Foreign Keys 69

Database Views 69
Database Users 70
Database Adapters 71
Storing Digital Signatures 72
Setting Up Your Database 72
Database Connections 73

Glossary 77

Index 81

Introduction

This *Administrator's Guide* provides information technology administrators need to successfully manage Netscape Application Server: Process Automation Edition (PAE). Read about the basics here before you begin configuring PAE. This guide assumes you have installed PAE on your system. For installation instructions, see the *Installation Guide*.

About This Guide

This guide provides information you need to use PAE:

- An introduction to PAE
- Several chapters providing detailed information about managing clusters and applications
- Information about using databases with PAE

Conventions Used in This Guide

File and directory paths are given in Windows format (with backslashes separating directory names). For Unix versions, the directory paths are the same, except slashes are used instead of backslashes to separate directories.

This guide uses URLs of the form:

http://*server.domain/path/file*.html

In these URLs, *server* is the name of server on which you run your application; *domain* is your Internet domain name; *path* is the directory structure on the server; and *file* is an individual filename. Italic items in URLs are placeholders.

This guide uses the following font conventions:

- The `monospace` font is used for sample code and code listings, API and language elements (such as function names and class names), file names, path names, directory names, and HTML tags.
- *Italic* type is used for book titles, emphasis, variables and placeholders, and words used in the literal sense.
- **Boldface** type is used for glossary terms

Viewing Documentation Online

For your convenience, PAE manuals are replicated online in both PDF and HTML formats. You can access the online documentation from the Help menu of each PAE component. You can access context-sensitive documentation by clicking a Help button or link in each PAE component.

Chapter

1

Introduction to PAE Administration

Netscape Application Server: Process Automation Edition (PAE) is a business process management system that runs on Netscape Application Server with a graphical design interface for defining forms and their routing, two administrative interfaces for configuring, and maintaining the PAE components, and an end-user interface for creating work requests, handling work items, and performing searches.

This guide focuses on using the information technology management interface provided by Process Administrator to manage clusters and applications. See the the *Business Manager's Guide* for information on the Process Business Manager interface.

This introductory chapter includes the following sections:

- Overview
    - Directories in PAE
    - Netscape Application Servers
    - PAE Applications
    - Security in PAE
- Process Administrator
- Directory Server Terms and Attributes

Overview

PAE consists of these components:

- Process Engine: The software internal to PAE.
- Process Administrator: An application accessible through your web browser that is used for managing clusters and applications in PAE.
- Process Business Manager: An application accessible through your web browser that is used for managing processes and statistics for PAE.
- Process Builder: The Java application for building PAE applications.
- Process Express: The HTML-based interface for end users who are accessing PAE applications.

PAE also uses these other components, which it associates into a **cluster**:

- a corporate user LDAP directory service
- a configuration LDAP directory service that stores the application definitions
- a relational database for user data, such as the products available from Oracle and Sybase
- one or more Netscape Application Servers
- a Netscape Enterprise Server
- a mail server for notifications

When an application developer using Process Builder deploys an application so that end users can access it, they must identify the cluster on which to deploy it. All successfully deployed applications are available to any valid end user on any Netscape Application Server across the cluster. All applications in a cluster share the same common database and directories. They access the same Directory Server for their process definitions and they use the same set of cross-application tables in the database, as well as the same corporate users and groups directory.

End users use Process Express to access the applications built in Process Builder. As they create new work requests and other examples of business processes, called **process instances,** and as they complete their assigned

tasks, called **activities** in the Process Builder and **work items** in Process Business Manager, they are generating user data that is stored in the cluster's database.

Figure 1.1 PAE Components

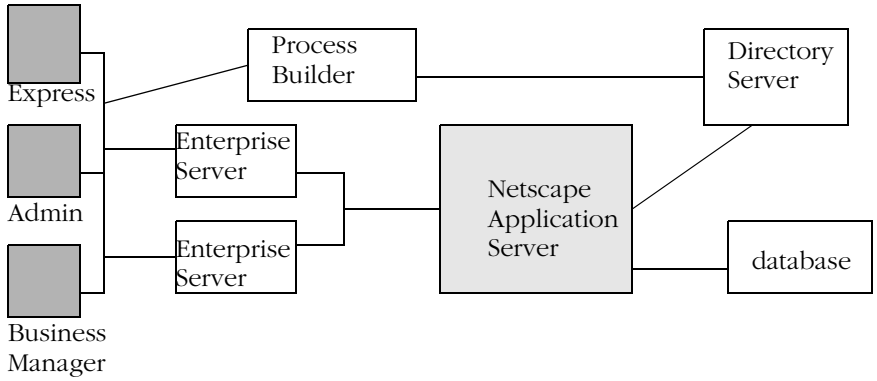

As the information technology administrator, you have these primary types of tasks:

- installing and configuring the software
- creating a cluster
- managing clusters and applications

Your first task is to install the PAE components and make sure they are configured correctly for your environment. See the *Installation Guide* for details.

The next step is to create a cluster so that Process Builder can deploy applications that can be used across the enterprise.

You can continue to manage and update the cluster as needed. For example, you may need to switch to a different corporate user directory or you may want to add other Netscape Application Servers to the cluster.

Directories in PAE

PAE uses directories for two purposes:

- to store PAE configuration information, such as process definitions
- to provide a list of the users and groups within a corporation who can be assigned to activities

The Configuration Directory

The configuration directory must be Netscape Directory Server 4.0. This directory stores PAE configuration information, including the application definitions for all deployed applications. The directory can also be used as a central repository for applications that are still being designed but are not yet deployed. The installation procedure extends this directory's schema to include the attributes and object classes required for PAE.

Once you define a cluster to use a particular configuration directory, you cannot switch to use another Directory Server in the cluster for your configuration information.

The Corporate User Directory

The corporate user directory must be Netscape Directory Server 4.0. This directory contains the set of corporate users who can be the assignee for a work item.

When you install PAE using all the defaults, you install a single Directory Server that you use for both types of information: users and configuration. This works well for using the sample applications and for initial testing of new applications where you can create a sample set of corporate users that you can test reliably.

Changing Your User Directory

If you want to do testing that simulates your production environment or if you want to move the application into production, you need to be able to access the actual users in your corporation. To do this, you need to change several default values to point to your company's corporate Directory Server, including the following:

- the directory that PAE uses for authenticating users and groups
- the access control rules (ACLs) for PAE-specific configuration styles
- your cluster's corporate directory
- make sure your sample applications have valid users

Note All Netscape Application Servers in a cluster must use the same information for all these settings.

To change the directory that PAE uses for authenticating users and groups, perform the following steps:

1. Go to the Server Administration page for your Netscape Enterprise Server. You access this page by going to the server:port defined for your 3.x SuiteSpot Administration Server. This URL is in the form *`yourServer`*:*`3.x Admin port`*.
2. Click the Global Settings button.
3. If not already chosen, click the LDAP Directory Server radio button to obtain directory service from an LDAP Directory Server.
4. Enter the correct information for your new corporate user directory:
   - hostname (defaults to your local machine)
   - port (defaults to 389)
   - Base DN (defaults to ou=People, o=mcom.com)
   - Bind DN and Bind DN Password are only required if you want to be able to add users or groups to the new directory.
5. Click Save Changes.
6. Restart the 3.x SuiteSpot Administration Server and all its servers.

This updates the `dbswitch.conf` file in the Enterprise Server. You can also manually edit this file instead of using the user interface to include the LDAP URL for the new directory. If you manually edit the file, put your new corporate directory as the second entry, which is the corporate directory URL.

To change the access control (ACL) for PAE styles, perform the following steps:

1. Go to the Server Administration page for your Enterprise Server.
2. Click the button for your Enterprise Server instance to go to its Server Manager pages.
3. Click the Configuration Styles button.
4. Click the Edit Style link.
5. Select Basic Auth from the drop-down list and click the "Edit this style" button.
6. Click the Restrict Access link. This displays the Access Control List Management page.
7. In the middle of the page, under "B. Pick an existing ACL," select Basic Auth and click the Edit Access Control button. This displays the Access Control user interface.
8. Click "anyone" in the first line. A second pane appears in the lower part of the page.
9. Click the "Authenticated people only" radio button in the lower pane.
10. Click the "All in the authentication database" radio button.
11. Make sure that the authentication method selected is Default.
12. Under Authentication Database at the bottom of the pane, select the new corporate directory from the drop-down list and click its radio button.
13. Click Update in this frame to update the ACL rules.
14. Click Submit in the top frame to set the change.
15. Click OK to save your changes.
16. Click Save and Apply to apply your changes to the server.
17. Repeat Steps 5-14 for the two other PAE styles: Admin IT Auth and Admin Data Auth.

There are two ways to make your applications use the new Directory Server:

- Make an existing cluster point at the new corporate directory.
- Create a new cluster that uses the new corporate directory.

If you are making an existing cluster point at the new directory, perform the following steps:

1. In Process Administrator, use the Change Cluster Information page to update the cluster with the new corporate user directory URL. See "Changing Cluster Information" in Chapter 3, "Clusters" for instructions.
2. Making sure you have access to the right directory from Process Builder. There are two ways of doing this:
    - If you have a cluster available during the design phase, you don't need to include the new corporate user directory's URL in the `preferences.ini` file. Instead, make sure your application uses the cluster's corporate directory. To do this, open the application's main properties inspector and set the Corporate Directory to be based on the cluster.
    - If you are designing an application without access to a cluster during the design phase, you need to add the new LDAP URL to the `preferences.ini` file. In this case, make sure your application uses a specific corporate directory. To do this, open the application's main properties inspector and pick the Corporate Directory you want to use. Note that if you deploy to a cluster that uses the same directory, the assignments work as designed.

If you want to create a new cluster, follow the following steps:

1. In Process Administrator, create a new cluster using the new directory. For more information, see "Creating a Cluster" in Chapter 3, "Clusters."
2. In Process Builder, redeploy your existing applications to the new cluster.

There are two ways to make your sample application users work with the new corporate user directory:

- Change the user assignments to point to users in the new directory.
- Add the required users to the new directory.

If the users and groups you use in your existing Process Builder applications also exist as valid users and groups in the new Directory Server, your applications will work as is.

If the original users and groups are not valid any more, you must change them so that they can be found in the new Directory Server and then redeploy the applications. Typically, you will need to replace the defaulted "admin" user with a valid user from the new corporate directory in the Users and Groups folder for each application. See "Sample Applications" in Chapter 10 of the *Developer's Guide* for details of other configurations you may need to set for specific applications.

If you want to change the user assignments, follow these steps:

1. Select a user or group in the application tree.
2. Open its properties inspector.
3. Pick a new user ID for the assignment.

If you want to add users or groups to the new directory, follow these steps:

1. Launch Netscape Console.
2. On the authentication dialog box, enter this information:
    - administrative user name
    - administrative password
    - administration URL for the new directory server's Administration Server, including the port number

3. Click on User and Groups tab.

4. Pick New User or New Group from the drop-down list in lower right corner and click Create.

5. Enter new user or group info and click OK when done.

Netscape Application Servers

Process Engine is an Enterprise Java Bean application that runs on Netscape Application Server. PAE uses the application server to run the HTML-based Process Administrator, Process Business Manager, and Process Express.

Each cluster must have at least one application server, but there can be more than one if there are several networked systems using the same cluster. All applications are replicated to all Netscape Application Server machines in a cluster.

As the administrator, you can perform these Netscape Application Server-related tasks:

- Join your local Netscape Application Server to an existing cluster.
- Remove your local Netscape Application Server from a cluster.

PAE Applications

All PAE applications run as applications on each Netscape Application Server machine in a cluster. Deployed applications are deployed to all application servers in a cluster, so if one server is unavailable, the application continues to run on the other machines in a cluster.

When a specific Netscape Application Server machine shuts down, all of the applications on that server also shut down. When the server comes back up, it automatically restarts all of its applications.

Security in PAE

PAE supports additional security features such as using SSL-enabled Enterprise Servers to provide secure content and access. PAE also allows designers to build applications that use certificates and digital signatures as part of their processing.

If you want to enable SSL on your Enterprise Server, read Chapter 8, "Using Encryption and SSL," in the online Enterprise Server Administrator's Guide. You can access the Enterprise Server help system by clicking any button on a Netscape server user interface form.

If you want to include digital signatures in a form, read the information about how to design with them in Chapter 6, "Defining Data Fields," in the *Developer's Guide*. Digital signatures are stored in a special database table, wf_blobs, so the administrator can query the database as needed to verify a signature. Also see "Storing Digital Signatures" on page 72 for more information.

For further information about security in general and about how to use the security features available in Netscape products, see the Security Documentation page in the DevEdge developer site, at
`http://developer.netscape.com/docs/manuals/security.htm`

Process Administrator

To perform administrative tasks, you must have the application server that you are using for PAE running on your local computer. Access to the Process Administrator's interface is through its home page at

`http://yourServer/Administrator.apm`

Process Administrator uses a tabbed HTML-based interface that provides access to management functions in these areas:

- Cluster management
- Applications

Cluster Management Forms

The Cluster Management tab displays different sets of forms depending on the situation: one set is for use in creating or joining a cluster and the other set is for managing existing clusters.

Create or Join a Cluster The first time you access the Cluster Management tab after installing PAE or the first time you access it from Netscape Application Server machine that is not already part of a cluster, the Create or Join a Cluster page is displayed.

Create Cluster Used to create a new cluster.

Join Cluster Used to join this application server to an existing cluster.

Cluster Management Used to manage the cluster that this application server belongs to.

Change Cluster Information Used to change information about an existing cluster.

Delete Cluster Used to delete a cluster.

Unjoin Cluster Remove this server from the cluster.

Application Forms

Process Administrator provides several management forms for applications. You can change the state of an application, check its logs, and archive and delete its data.

Directory Server Terms and Attributes

Because much of PAE is dependent on Directory Servers, this section is included to help clarify some of the most relevant concepts and terminology.

Whether you are accessing the corporate user directory for your set of users or defining a cluster in the configuration directory, you need to understand how to identify the directory and the specific cluster entry within the server.

There are some standard LDAP terms and attributes that you may need to understand before you can create Directory Server entries. This section briefly describes them for your convenience. For detailed information, see the Directory Server manuals, which you can access by clicking on any help button in a Directory Server product.

LDAP Terms

In general, Netscape Directory Servers use standard LDAP terminology, but different administrative forms may use slightly different sets of equivalent terms. Common terminology you may encounter as you install and use the Directory Server includes the following:

Distinguished Name (DN) A series of comma-delimited attributes that uniquely identify the directory entry location within the directory tree. This could be a person, a group, an organization, or any other object for which you want to maintain information in a directory. In the case of PAE, information about a cluster is maintained in a directory.

Base DN The entry at which to start directory searches, sometimes referred to as the *search base*. This base is often the root entry, that is, the search starts at the top of the directory tree.

Bind DN The DN used to access the directory. Directory Server authentication is referred to as *binding* to the directory. Which DN you use as the Bind DN determines the level of directory access permitted. This is often the root DN, who has complete access to the directory, and so the Bind DN is sometimes referred to as the *unrestricted user*. The default Bind DN for Netscape Directory Servers is `cn=Directory Manager`.

Directory Suffix A distinguished name (DN) suffix for your local directory. All incoming LDAP queries must contain this suffix, which is equivalent to the root entry of your Directory Server structure. This provides the highest level of identification for a specific directory. For example, `o=airius.com`. Everything contained within a directory is underneath this entry. If you know the directory suffix or root entry for a directory, you know which directory it is.

Root Entry The first entry in a directory tree, that is the top of the tree. This is often, but not always the Base DN. The root entry corresponds to the directory suffix. If you know the root entry or directory suffix for a directory, you know which directory it is.

LDAP Attributes

When you identify a directory's location in a Directory Server's tree, such as when you define a cluster within the configuration directory, you typically need to use only a small set of LDAP attributes. These include the following:

- c (country)
- o (organization)
- ou (organizational unit)
- cn (common name)
- uid (user ID)

The common name entry of `cn=Directory Manager` is the default administrative user identifier for Directory Servers. It is set when you perform a default installation of the Directory Server.

Note You cannot use commas within an attribute value, only as delimiters between attributes.

Figure 1.2 A sample corporate user directory structure

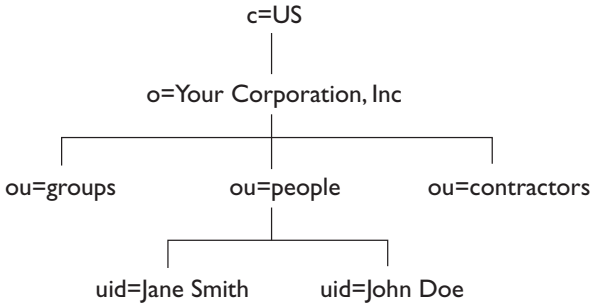

Directory Structure

The Directory Server uses a tree structure to define different sets of data. In a simple case, such as identifying a cluster, you could have a structure like this:

- `o=airius.com` (tree root level)
- `cn=My Cluster` (specific cluster branch)

If you had another cluster in the tree, you could have these values:

- `o=airius.com` (the same tree root level)
- `cn=Your Cluster` (a different branch of the Directory Server tree)

Together these two values uniquely identify the location of the cluster's directory entry in the Directory Server and are referred to as the cluster's *distinguished name*, or *DN*. When you want to uniquely identify the cluster, you need to include the entire DN, with attributes separated by commas and listed in order from most specific to highest level. For example, for a cluster, you could use this DN:

```
cn=My Cluster, o=airius.com
```

Figure 1.3 A sample directory containing a cluster

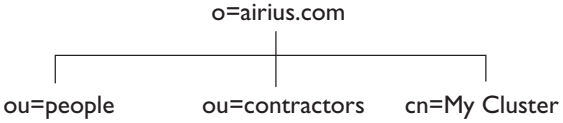

LDAP URLs

You use these attributes to identify the corporate user directory and your cluster entry to the Process Builder in the `preferences.ini` file after creating a cluster. This file requires you to use a specific LDAP URL format when you enter this information.

The Corporate Directory URL

If you do not require user authentication, as is typical for the corporate user directory, use this format:

```
ldap://yourDirServer:port/Base DN
```

For example:

```
ldap://netscape.mcom.com:389/o=mcom.com
```

The Cluster URL

If you require user authentication, as you do for the cluster entry, use this format:

```
ldap://Bind DN:Bind Password@yourDirServer:port/cluster DN
```

For example:

```
ldap://cn=Directory Manager:netscape@netscape.airius.com:389/cn=HR Cluster, o=mcom.com
```

Chapter

2

Configuration Files

Netscape Application Server: Process Automation Edition (PAE) uses several configuration files to define installation values, server configurations, clusters, and application versions.

Much of what they do is behind the scenes, but you may need to know about them from time to time. This chapter describes what you need to understand about each and how you might need to modify them.

This chapter describes these topics:

- Overview
- PAE Files and Folders
- Error Messages and Logs

Overview

Configuration files are included in the components that you install as part of PAE. The configuration files that are important for PAE administrators touch on several areas:

- The Enterprise Server Files
- The Process Builder Preferences File

The Enterprise Server Files

The Enterprise Server has several configuration files, but the three that are relevant to PAE are the Enterprise Server object configuration file (`obj.conf`), the `dbswitch.conf` file, and the `generated.https-`*`yourServerName`*`.acl` file.

- `obj.conf`: Object configuration values for the Enterprise Server, including access control information. These lines are added to your `obj.conf` as part of the PAE installation:

  ```
  NameTrans fn="assign-name" name="Admin Data Auth" from="/*Data.apm"
  NameTrans fn="assign-name" name="Admin IT Auth" from="/*IT.apm"
  NameTrans fn="assign-name" name="Basic Auth" from="/*.npm"
  ```

  If you specified a particular PAE administrator during installation, that user is assigned to the Admin Data Auth and Admin IT Auth groups with the same access privileges.

- `dbswitch.conf`: Contains an LDAP URL about the corporate directory that PAE uses to authenticate users.

  The following is a sample LDAP URL:

  ```
  ldap://xyz.mcom.com:389/ou=People, o=mcom.com
  ```

  This URL is obtained from the user during installation.

- `generated.https-`*`yourServerName`*`.acl`: Contains access control list (ACL) information about PAE.

The following lines are added as part of your PAE installation:

```
acl “Admin Data Auth”;
authenticate (user,group) {
prompt = “Netscape Process Manager Administration Data Manager”;
database = “CorporateDirectoryYOURDIR”;
};
deny (all) (user = “anyone”);
allow absolute (all)(user = “admin”);
```

Note that the default administrative user is `admin`, but users can specify any administrative user at installation.

The Process Builder Preferences File

The Process Builder has a `preferences.ini` file that identifies which cluster and corporate directory to use when designing applications.

You can add multiple cluster and directory entries to offer a choice to designers as they build applications in Process Builder. Each application can only belong to one cluster at a time and have access to a single user directory at a time, but a designer may work with several clusters at a time.

The `preferences.ini` file provides the unique cluster DN and identifies which corporate user directory is included in the cluster. If you cut and paste the two lines that result from a successful cluster creation, you are providing the `preferences.ini` file with one cluster and one corporate user directory. See “Creating a Cluster and Getting Started” in the *Installation Guide* for information about this hand-off to Process Builder. See “Directory Server Terms and Attributes” in Chapter 1, “Introduction to PAE Administration” for details about how to structure LDAP information.

Note If the process designer and the administrator are not the same individual, the administrator must give the following URLs to the process designer for inclusion in the `preferences.ini` file.

- To add a cluster entry, use this LDAP URL format:

`cluster=ldap://`*`Bind DN`*`:`*`Password@yourConfigDirServer`*`:`*`port/clusterDN`*

- To add an entry for a corporate user Directory, use this LDAP URL format:

`corp_dir=ldap://`*`corporateDirectoryServer`*`:`*`port/corpDirectoryBaseDN`*

PAE Files and Folders

The PAE installer adds many files and folders to your local system as it installs the PAE components.

The PAE files are added to *serverRoot*/bpm.

Process Builder files and folders are added to the *serverRoot*/builder folder.

Error Messages and Logs

There are several logs in PAE: some are for Process Administrator errors or information, others are specific to an application. You can directly open an application's logs or the PAE logs in your web browser window.

- Process Builder logs are located at: *serverRoot*/builder/support/log/error.html and warning.html.
- Process Administrator logs are located in the *serverRoot*/bpm/admin/logs folder. They are info.html, error.html, and security.html.
- PAE applications use several log files, with each Netscape Application Server in a cluster having its own set of log files. PAE administrators can view these application logs that are maintained for each application on a per-server basis:
  - information log (info.html): shows information about the application's operations.
  - error log (error.html): shows the errors that have occurred with the application.
  - security log (security.html): shows the security violations that have occurred with the application.

  The application log files are located at *serverRoot*/bpm/applications/*applicationName*/logs.

There are other log files for each of the Netscape servers:

- Enterprise Server: *serverRoot*/https-*yourEnterpriseServer*/logs
- Administration Server: *serverRoot*/admin-serv/logs
- Directory Server: *serverRoot*/slapd-*yourDirServer*/logs

Chapter

3

Clusters

Netscape Application Server: Process Automation Edition (PAE) associates a configuration directory, a corporate user directory, a relational database, an Enterprise Server and one or more Netscape Application Servers into a **cluster**. Applications are deployed to a cluster and are copied to all Netscape Application Servers in the cluster. All applications in a cluster access the same set of shared components: database, directories, and Netscape Application Servers.

Each Netscape Application Server that is part of PAE is associated with one and only one particular cluster. When designers deploy a new application, they must identify which cluster it belongs to. PAE replicates all applications across all Netscape Application Servers in a cluster and the administrator can manage all applications from any Process Administrator in the cluster.

This chapter describes these topics:

- About Clusters
- Creating a Cluster
- Joining an Existing Cluster
- Managing a Cluster

About Clusters

A cluster associates these components:

- an LDAP-compliant corporate user directory
- a relational database
- an LDAP-compliant configuration directory
- one or more Netscape Application Servers, each with its own Process Administrator

The initial cluster includes one Netscape Application Server. Later you can join additional Netscape Application Servers to an existing cluster as needed. Note that each Netscape Application Server can only belong to one cluster, and you must have a cluster available in order to use PAE.

Each PAE-enabled Netscape Application Server has its Process Administrator. You can use any Process Administrator and get the same view of a cluster because all applications are fully replicated to all Netscape Application Servers in a cluster.

As the PAE administrator, you can perform these tasks that relate to clusters:

- Create a new cluster.
- Join your Netscape Application Server to an existing cluster.
- Change cluster information.
- View the Process Administrator logs.
- Remove your Netscape Application Server from a cluster.
- Delete a cluster.

Creating a Cluster

You administer PAE through the Process Administrator interface on an Enterprise Server. The first time you access Process Administrator's home page after installing PAE, Process Administrator detects that this Netscape Application Server is not yet associated with any cluster and immediately prompts you to create or join a cluster.

Note The database user who is creating the cluster and all the other database users for applications (if any) should be ordinary users with create table/view permissions. None of them should be the dbo of the database.

To create a cluster, follow these steps:

1. Launch your web browser.
2. Go to your local Process Administrator's home page at `http://yourServer/Administrator.apm`

   This initializes Process Administrator and displays the Create or Join Cluster page. This step may take longer the first time you initialize Process Administrator after installing or restarting your server.

Figure 3.1 The Create or Join Cluster page

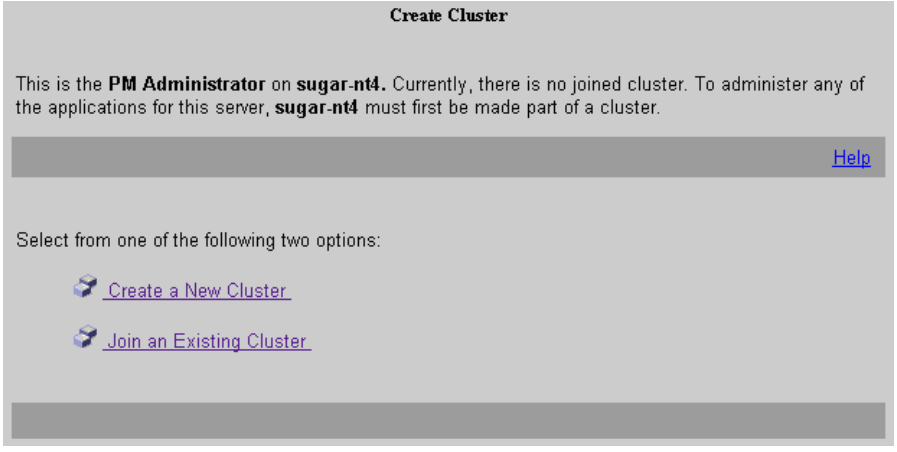

3. Click the "Create a New Cluster" link. This displays the Create Cluster page.
4. Enter information about these components (For details, see the next section, "New Cluster Information"):

- cluster
- corporate users and groups directory
- configuration directory
- database information
- mail server information

5. Click Create Cluster when you are done.

   A series of diagnostic messages are displayed onscreen as the cluster is created. The last few lines contain instructions for modifying the Process Builder `preferences.ini` file.

6. Copy the lines from the onscreen messages and paste them into the `preferences.ini` file. This file identifies to the Process Builder which Directory Server to use as the corporate user directory and where one or more clusters are located in the configuration directory. You need to do this before a designer can access this cluster from Process Builder.

   The Process Builder's `preferences.ini` file is located in the Process Builder root folder. This is the *`serverRoot`*`/builder` folder by default.

Once you've completed these steps, designers can deploy applications to a cluster.

New Cluster Information

The administrator must provide information about a new cluster. The only general restriction is that you cannot use double quotation marks (").

Figure 3.2 The Create Cluster page

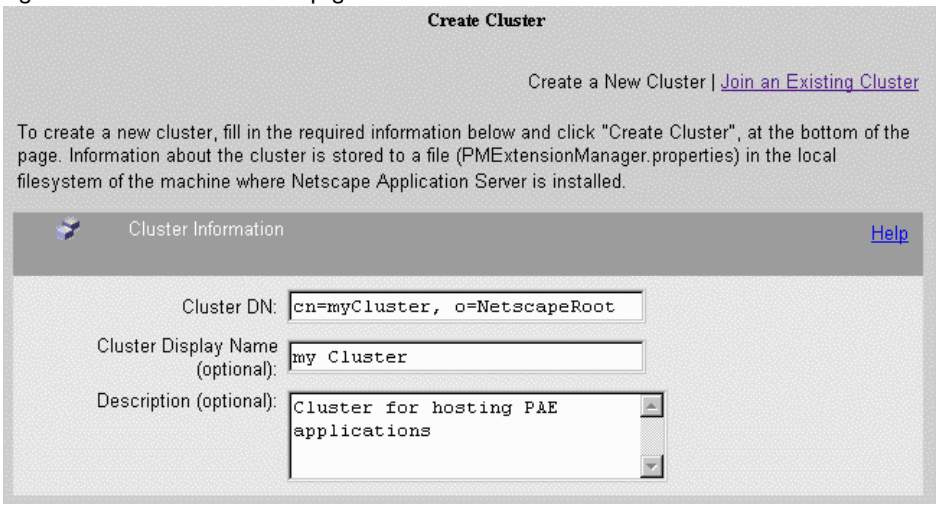

General Cluster Information

Cluster DN The full distinguished name for your cluster. That is, the directory path for the cluster data. For example, `cn=PM Cluster, o=NetscapeRoot`. For more information about what LDAP attributes to enter here, see "Directory Server Terms and Attributes" in Chapter 1, "Introduction to PAE Administration."

Cluster Display Name The display name you choose for your cluster. (Optional)

Description A more meaningful description of the cluster. (Optional)

Figure 3.3 The directory information sections

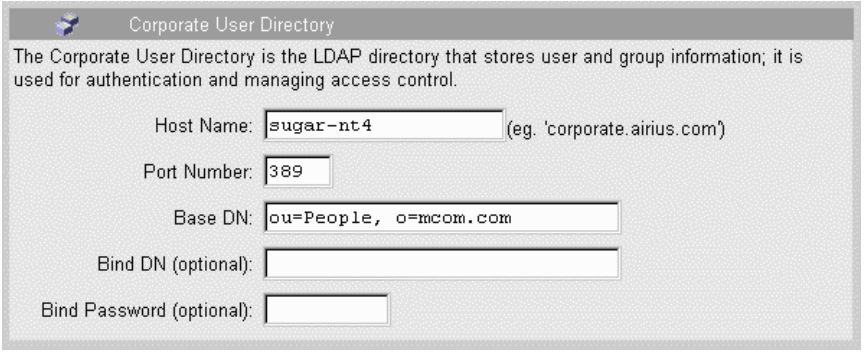

For further information on Directory Server terminology, see "LDAP Terms" in Chapter 1, "Introduction to PAE Administration."

Corporate User Directory Information

Host Name The full host name for your corporate user directory. For example, `corporate.airius.com`.

Port Number The port number for this directory. This defaults to 389 if you are using a Netscape Directory Server.

Base DN The base distinguished name for the directory. For example, `o=mcom.com`. For more information about what LDAP attributes to enter here, see "Directory Server Terms and Attributes" in Chapter 1, "Introduction to PAE Administration."

Bind DN (Optional) A valid user ID for a user of this directory. Leave this blank if PAE will be accessing the corporate directory as an anonymous user, which is the default. This allows read-only access.

Bind Password (Optional) The password for this directory user. Leave this blank if PAE will be accessing the corporate directory as an anonymous user.

Figure 3.4 The configuration directory information section

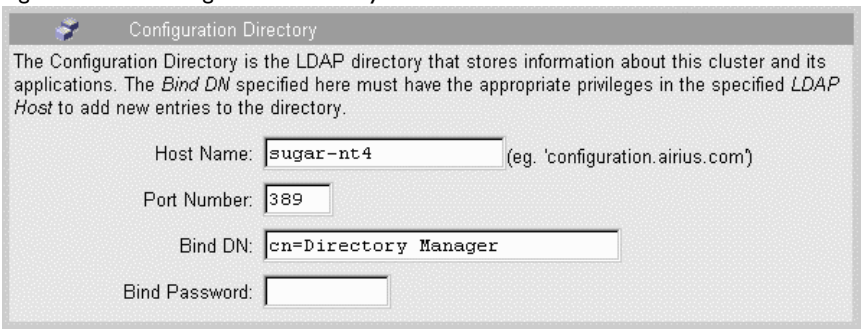

Configuration Directory Information

Host Name The full host name for the Directory Server that you plan to use for your cluster and application configuration. For example, `netscape.mcom.com`.

Port Number The port number for this directory, which defaults to 389 because this is a Netscape Directory Server.

Bind DN A valid user ID for a user of this directory. This user must have full Directory Server privileges. For example, the installation default is `cn=Directory Manager`. For more information about what LDAP attributes to enter here, see "Directory Server Terms and Attributes" in Chapter 1, "Introduction to PAE Administration."

Bind Password The password for this directory user.

Figure 3.5 The database information section

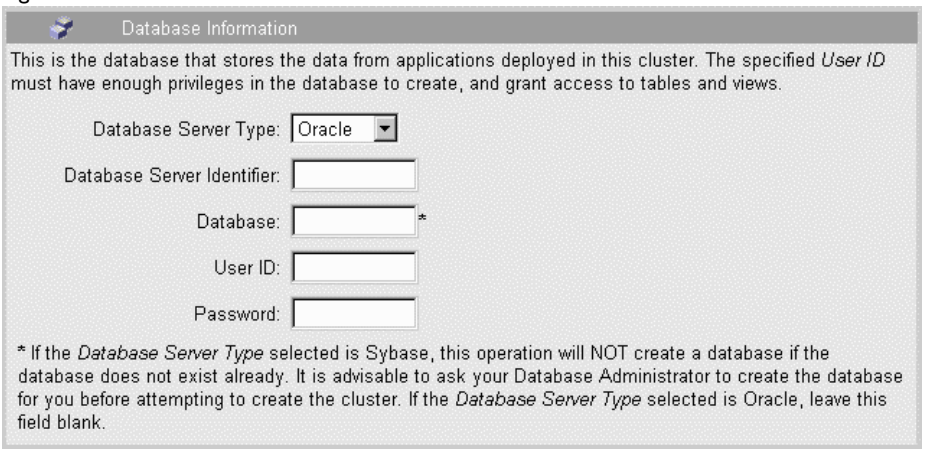

Database Information

Database Server Type The type of database you are using from the drop-down list. The choices are Oracle and Sybase.

Database Server Identifier The database server name that you used when you installed the database. See Chapter 5, “Databases” for more information about databases.

Database For Oracle, leave this blank. For Sybase, use the name for the database that Process Administrator should use for creating tables and views. Note: On Sybase, this database must already exist.

User ID The user ID for a valid user for this database. The user must have enough privileges to create tables and write to them.

Password The password for this database user.

Note If you are using Sybase on Windows NT, be sure that your DSQuery system environment is set to point to the correct sybase database server. If not, you will not be able to connect to the database.

Figure 3.6 The mail server information section

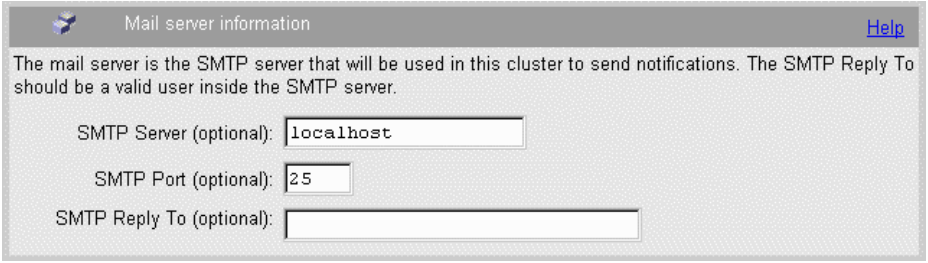

Mail Server Information

SMTP Server (Optional) The mail server associated with the cluster.

SMTP Port (Optional) The mail server port.

SMTP Reply To (Optional) The reply-to address that appears on notifications sent to users by PAE applications. This address could be set to the PAE administrator's email address, so that the administrator would be the one handling any notification replies users send.

What Happens at Cluster Creation

When you create a new cluster, Process Administrator performs these operations:

- creates entries in the configuration directory
- creates database tables
- creates a cluster definition file (`PMExtensionManager.properties`) in your local PAE folder, at *`serverRoot`*`/bpm`. This file stores the machine name, port number, bind DN, and bind password for the configuration directory.
- displays the URLs for the cluster and the corporate user directory so you or the designer can paste these lines into the `preferences.ini` file.

Joining an Existing Cluster

The initial cluster includes one Netscape Application Server, but you can add additional servers at any time. Whenever you access Process Administrator for the first time from a new Netscape Application Server, Process Administrator checks which cluster that server belongs to. If it is not associated with a cluster, you are prompted to create a cluster or join one of the clusters that already exists.

To join an existing cluster, follow these steps:

1. Go to your local Process Administrator's home page at `http://yourServer/Administrator.apm`. This displays the Create or Join Cluster page.
2. Click the "Join an Existing Cluster" link. This displays the Join Cluster page.
3. Enter information about these components (For details, see the next section, "Existing Cluster Information"):
   - the cluster you want to join
   - the cluster's configuration directory
4. Click Join Cluster when you are done.

This operation creates a cluster definition file in the PAE folder on your local machine, at

`serverRoot/bpm/PMExtensionManager.properties`

Existing Cluster Information

The administrator must provide this information about an existing cluster. The only general restriction is that you cannot use double quotation marks (").

Figure 3.7 The Join Cluster page

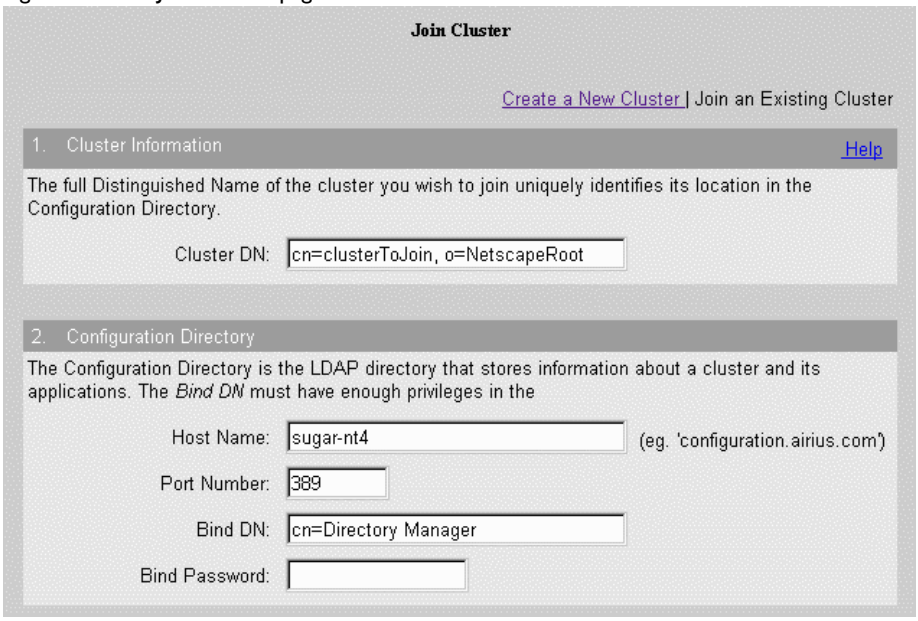

General Cluster Information

Cluster DN The full distinguished name for your cluster. This is the directory path for the cluster data. For example, you could enter `cn=PM Cluster, o=airius.com`. For more information about what LDAP attributes to enter here, see "Directory Server Terms and Attributes" in Chapter 1, "Introduction to PAE Administration."

Configuration Directory Information

Host Name The full host name for the Directory Server you plan to use for your cluster and application configuration. For example, `sample.airius.com`.

Port Number The port number for this directory, which defaults to 389 because this is a Netscape Directory Server.

Bind DN A valid user ID for a user of this directory. This user must have full Directory Server privileges. For example, you could enter `cn=Directory Manager`. For more information about what LDAP attributes to enter here, see "Directory Server Terms and Attributes" in Chapter 1, "Introduction to PAE Administration."

Bind Password The password for this directory user.

What Happens at Cluster Joining

When you join an existing cluster, Process Administrator performs these operations:

- creates a file, `PMExtensionManager.properties`, on the Netscape Application Server machine
- registers one AppLogic for each application in the existing cluster with Netscape Application Server

Managing a Cluster

Process Administrator provides a main Cluster Management page at the Cluster Management tab.

This section includes these topics:

- Changing Cluster Information
- Accessing Directory Server Information
- Viewing the Process Administrator Logs
- Unjoining From a Cluster
- Deleting a Cluster

Figure 3.8 The Cluster Management page

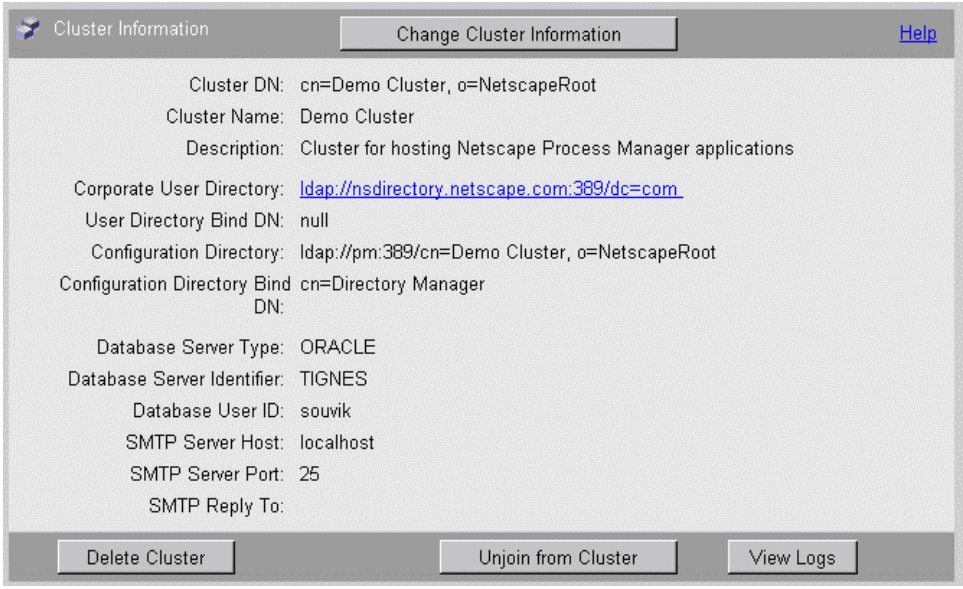

Changing Cluster Information

As server administrator, you can change some cluster information. You can change the cluster's name and description and you can change which corporate user directory you include in the cluster You cannot switch to use a different configuration directory or database in this cluster.

Note If you change the corporate user directory, you may wish to update the Process Builder `preferences.ini` file with its LDAP URL. When Process Builder designers want to revise a deployed application, if they define the application's corporate directory as cluster-based, then the new user directory is automatically used. Designers need to check that their group and role assignments are still valid whenever the corporate user directory changes.

You must also change the corporate directory on your Enterprise Server to match that of the LDAP URL. You can change this in the `dbswitch.conf` file.

To change a cluster's information, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Cluster Management tab.
3. Click the Change Cluster Information button. This displays the Change Cluster page.
4. You can enter information about these cluster components (For details, see the next section, "New Cluster Information"):
   - cluster (name, description, and directory server information)
   - database (password only)
5. Click Apply Changes when you are done.

Figure 3.9 The Change Cluster Information page

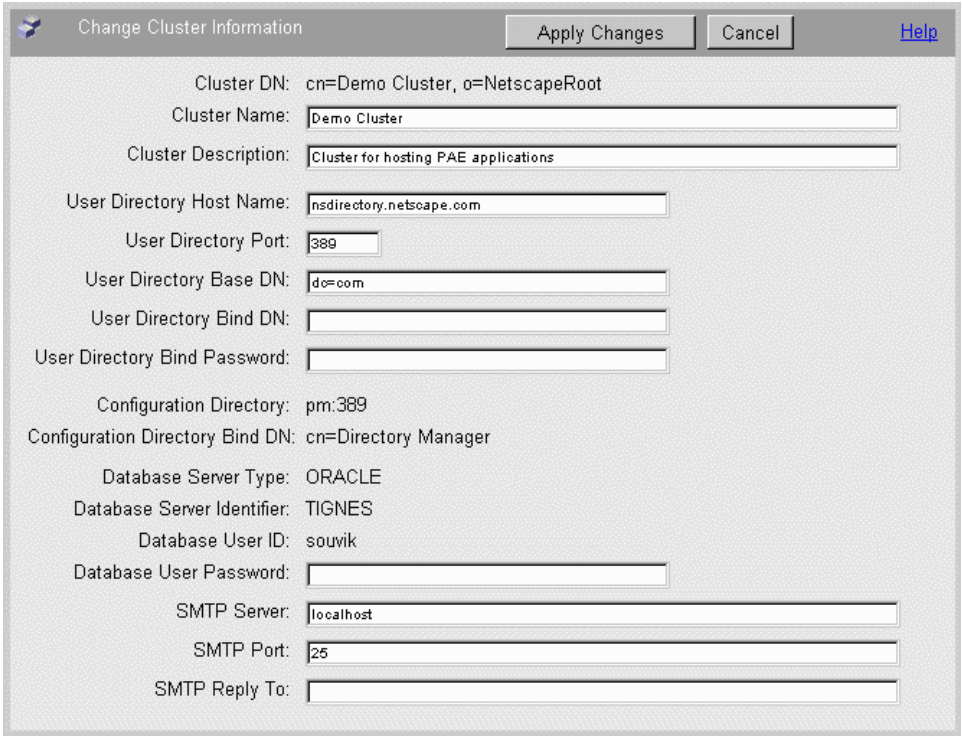

New Cluster Information

The administrator can modify the following information about an existing cluster. The only general restriction is that you cannot use double quotation marks (").

Cluster Name The user-defined display name for the cluster.

Cluster Description Text describing the cluster.

User Directory Host Name The server name for the corporate user directory.

User Directory Port The port number for the corporate user directory. This defaults to 389 for Netscape Directory Servers.

User Directory Base DN The directory suffix (or "base DN") for the corporate user directory. For more information about what LDAP attributes to enter here, see "Directory Server Terms and Attributes" in Chapter 1, "Introduction to PAE Administration."

User Directory Bind DN (Optional) The user ID for a user of this directory server. Leave this blank if PAE will be accessing the corporate directory as an anonymous user.

User Directory Bind Password (Optional) The password for this directory user. Leave this blank if PAE will be accessing the corporate directory as an anonymous user.

Database Password The database user's password.

SMTP Server Host The mail server associated with the cluster.

SMTP Server Port The mail server port.

SMTP Reply To The reply-to address that appears on notifications sent to users by PAE applications. This address could be set to the PAE administrator's email address, so that the administrator would be the one handling any notification replies users send.

Accessing Directory Server Information

You can access Directory Server information by clicking one of the LDAP directory links on the Cluster Management page.

Viewing the Process Administrator Logs

You can view the Process Administrator logs by clicking the View Logs button.

Figure 3.10A sample Process Administrator's log

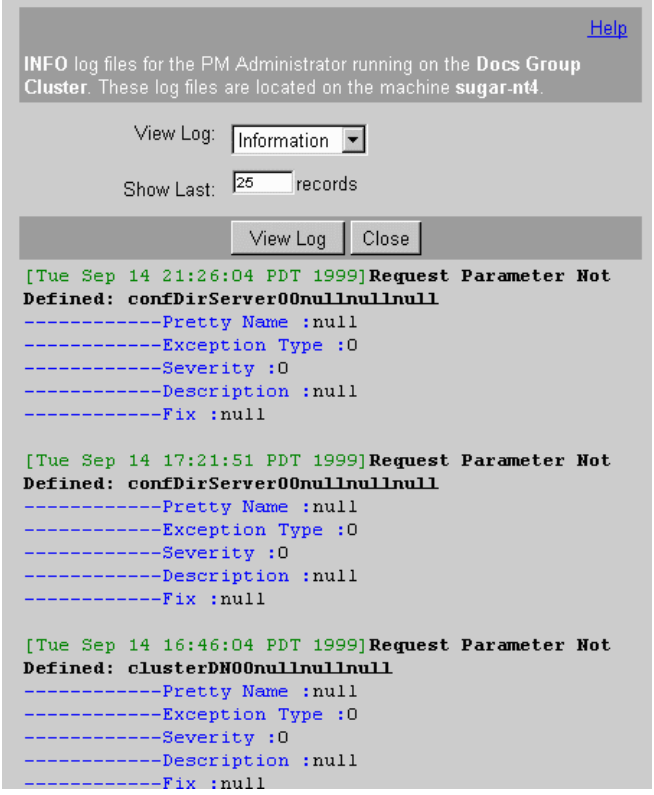

If Process Administrator is not available, you can access the Process Administrator logs (`error.html`, `info.html`, `and security.html`) at this location:

serverRoot`/bpm/admin/logs`

Unjoining From a Cluster

You can unjoin an application server from a cluster. This revises the entries in the configuration directory and removes the applications that are stored on that server machine.

To unjoin an application server from a cluster, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Cluster Management tab.
3. Scroll down on the Cluster Information page.
4. Click the "Unjoin from Cluster" button.

After you unjoin the server from its cluster, the next time you access Process Administrator, you will be prompted to create or join a cluster.

When you unjoin a server from a cluster, Process Administrator performs these operations:

- removes that server's entry from the configuration directory
- deletes the `PMExtensionManager.properties` file from the server's file system

Deleting a Cluster

You can delete a cluster that you no longer need. When you delete a cluster, Process Administrator deletes the cluster entry in the configuration directory and deletes the local application server copy of the `PMExtensionManager.properties` file. You can also choose to delete the database tables.

Note When you delete a cluster, you need to inform designers using Process Builder that the cluster's entry in the Process Builder `preferences.ini` file is no longer valid. If they do not not want to delete the references to the cluster in the `preferences.ini` file, the cluster appears in their Select An Application window with a red X through it, indicating that it is unavailable.

To delete a cluster, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Cluster Management tab.
3. Scroll down on the Cluster Information page.
4. Click Delete Cluster. This displays the Delete Cluster page.

Note You can choose to delete the application data for all cluster applications, which also deletes the application-specific database tables. This is not recoverable. If you are not sure whether you want to do this, even for an old unused cluster, first archive the application data. For details, see the section "Exporting and Deleting Data" in Chapter 4, "Applications."

Figure 3.11 The Delete a Cluster page

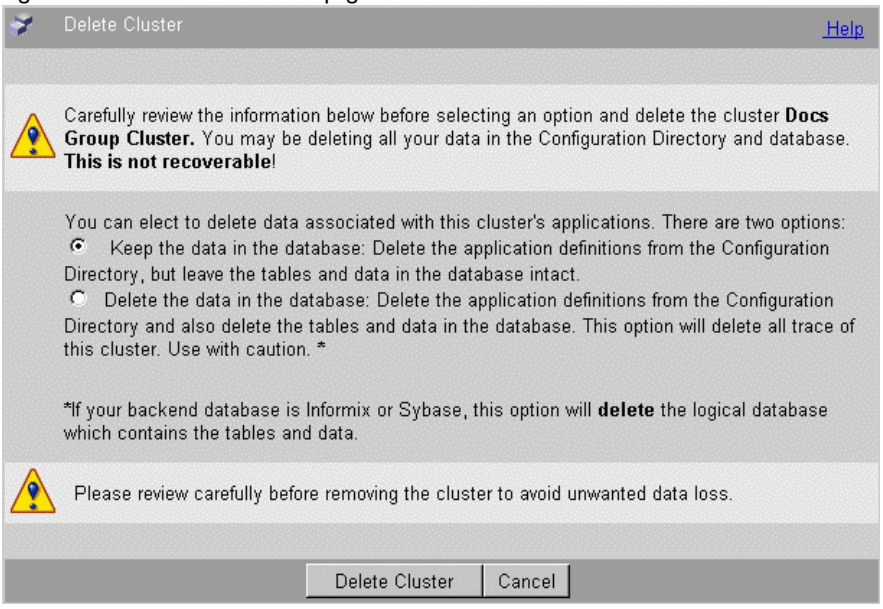

5. If you want to keep the application data, select the "Keep the data in the database" radio button.

   If you want to delete the application data, select the "Delete the data in the database" radio button.

6. Click Delete Cluster to delete the cluster (and your user data if you chose that option).

If you try to delete a cluster that still has other active servers running, Process Administrator first removes the other servers, then it removes your local server and deletes the cluster. Note that this also removes the `PMExtensionManager.properties` file from your local server, so the next time you try to access Process Administrator, you are "initializing" it again and you will have to create a new cluster.

Chapter

4

Applications

Applications that handle business processes are at the heart of the Netscape Application Server: Process Automation Edition (PAE) product. PAE provides several forms for managing applications. You can change an application's state, view its logs, and export and delete its data.

This chapter describes these topics:

- About Applications
- Stopping & Starting Applications
- Closing Applications
- Uninstalling Applications
- Viewing the Application Logs
- Exporting and Deleting Data

About Applications

Designers use Process Builder to define a PAE application. When the application is complete, the designer can deploy it to a cluster for final testing or for normal production use by Process Express end users.

A deployed **application** consists of a process definition that is stored in the configuration directory and a set of user data stored in the cluster's database.

Once an application is deployed, Process Administrator maintains information about its state. As the administrator, you can change this information, and as you do so, Process Administrator keeps the application information up to date in the configuration directory.

Figure 4.1 The Deployed Applications page

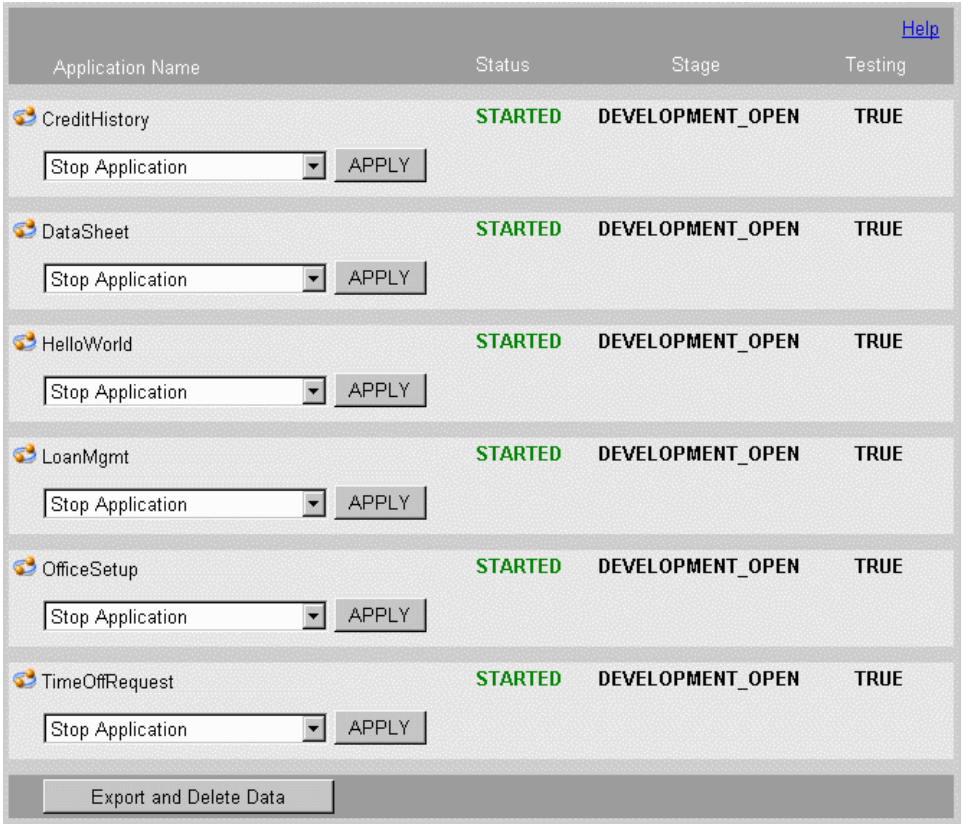

Application states include both their status and their stage. An application's **status** indicates whether it is currently running or not. An application's **stage** indicates whether it is in development or production, and whether it is available to end users.

Application statuses are:

- Started (On): End users and administrators can access the application to initiate and manage process instances and work items.
- Stopped (Off): End users cannot initiate new process instances for the application and cannot execute actions on existing process instances.

Administrators can stop started applications or start a stopped application.

Designers deploy to Development or Production. If they deploy to Development, they can still change almost all aspects of the application. If they deploy to Production, they are very limited in what changes they can make.

Application stages are:

- Open: This is the initial stage after a designer has deployed the application.
- Closed: The administrator has closed down the application. End users can complete existing process instances but cannot initiate any new process instances. The PAE administrator can reopen closed applications at any time.
- Obsolete: The administrator has uninstalled a closed application, and may have removed the application's user data. End users can no longer work on process instances for this application. Only closed applications can be uninstalled. Once an application has been uninstalled, you cannot resurrect it.

Another stage, Design, is only visible and accessible from Process Builder. This stage means that the application is not yet fully deployed, but its definition is stored in the configuration directory to provide a centralized repository so that multiple designers can always obtain the most current standard version of the application definition. The designer puts an application into this stage by deploying it as "Stored Only" from the Deploy Applications dialog box.

Finally, an application can also be in testing mode. The designer has deployed the application, but all work items are set as assigned to their initiator (which is typically the designer or the administrator) to permit further testing of all aspects of the process definition. Once an application is deployed to production, it can no longer return to the testing stage.

When the administrator changes an application's stage, Process Administrator updates the application definition in the configuration directory.

Note If a designer deploys a locally saved version of an already deployed application, and the two applications do not have the same stage value, the local version does not overwrite the stage of the application already in the cluster.

Stopping & Starting Applications

Among the actions that an administrator can perform on an application are stopping and starting it. These change its status, which is maintained in the configuration directory as part of the persistent application data that describes all instances of the application across the cluster.

Figure 4.2 Application administration actions

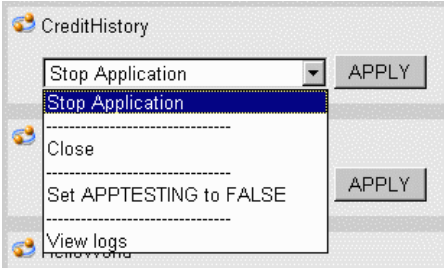

If an application is started and the administrator stops it, this shuts it down everywhere across the cluster. This sets the application's status to Stopped, which does not allow access to process instance detail by end users

If an application is stopped and the administrator starts it, this sets the application's status to Started, which allows full access to process instances and work items by end users and administrators.

When it starts an application, either as part of an autostart or as part of a manual start, Process Administrator performs these operations:

- Retrieves the application's definition from the configuration directory.
- Establishes a connection with the database.

Autostart

Whenever an end user accesses an application in the Process Express, PAE checks the application's status. If the application's status is Started but it isn't running, PAE does an autostart. This can occur in two situations:

- The application has just been newly deployed.
- The server through which an end user is accessing the application was unavailable and has now been restarted.

Because Process Administrator has to perform several operations when it starts an application, there can be a time delay for the first end user who accesses an application in these situations. As a result, some administrators may prefer to force initializing of an application immediately after it is newly deployed or to force initialization of all applications after restarting a server.

Initializing an Application

An application is initialized when someone such as an end user or the administrator invokes it as a web browser URL. As part of the initialization operation, an application checks the configuration directory to make sure it has the latest definition. If the configuration directory's definition is more up to date, the application downloads that definition to server.

To initialize an application, use this URL:

```
http://malta.mcom.com:4080/cgi-bin/gx.cgi/AppLogic+ContentStoreTest.npm
?eventId=OnListEntryNodes
```

The result is that the application is already up and running when the first user accesses it. To initialize all the applications on a server requires you to repeat this type of URL for each application. You may choose to only initialize the most time-critical applications.

Stop

You can use Process Administrator on a given Netscape Application Server to stop an application. For example, you may want to stop an application when you do not want end users to be able to access it. Stopping an application sets its status to Stopped in Process Administrator.

If an administrator stops an application, users are prevented from creating new process instances, work items, or otherwise accessing that application. Database transactions running when the administrator stops the application are allowed to complete.

To stop an application, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Applications tab. This displays the list of applications in the cluster.
3. For the application you want to stop, choose the Stop Application command from the drop-down list for that application.
4. Click Apply.

Now you and your end users cannot access the process instance detail for the application. End users cannot initiate new process instances or continue working on existing work items.

Start

You can start a stopped application. When you start an application, you are setting its status to Started in Process Administrator.

To start an application, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Applications tab. This displays the list of applications in the cluster.

3. For the application you want to start, choose the Start Application command from the drop-down list for that application.

4. Click Apply.

Now you and your end users can access the process instance detail for the application from any server in the cluster. End users can once again initiate new process instances and continue working on existing work items.

Closing Applications

Closing an application means that end users cannot initiate any new process instances, but they can continue to complete existing process instances. Closing an application sets its stage to Closed.

When you no longer need an application, you can close it. When you uninstall an application, you cannot "reinstall" it.

End users can continue to work on work items until there are no open work items for the application, at which time you can uninstall the application.

To close an application, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.

2. Click the Applications tab. This displays the list of applications in the cluster.

3. For the application you want to close, choose the Close Application command from the drop-down list for that application.

4. Click Apply.

End users cannot initiate new process instances but they can continue working on existing work items.

Uninstalling Applications

If you have an application that you no longer need, you can manually uninstall it. The application must already be in a closed stage without any remaining open process instances.

Typically, a PAE administrator follows this sequence of operations:

1. Leaves an application in the closed stage for a while.
2. Manually terminates any open process instances that do not complete within a certain period of time.
3. Exports the application's process instances (that is, the user data) and may choose to delete it from the database.
4. Uninstalls the application from the cluster.

Note If you are planning to delete a cluster, it is strongly recommended that you first uninstall each of its applications to avoid additional manual clean-up tasks after you delete the cluster.

Occasionally, you may have applications that you no longer want to use, and that you want to remove from the end user interface. You can uninstall a closed application, which marks it as Obsolete. Designers can delete obsolete applications from the Process Builder, but they cannot redeploy them. When you uninstall an application, you stop end users from initiating any new process instances from the New Process page.

When you choose the Uninstall command, you are asked if you want to delete the application-specific database tables. If you confirm the deletion, the application's user data is removed. If not, the data is left intact.

You cannot uninstall an application that has any open process instances. You need change the state of any open process instances to terminated if you want to continue with uninstalling the application.

To uninstall a closed application, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Applications tab. This displays the list of applications in the cluster.
3. For the application you want to uninstall, choose the Uninstall Application command from the drop-down list for that application.
4. Click Apply.

   This displays the Uninstall Application dialog box.
5. Check the radio button to leave the application's database table intact or to delete it, depending on your situation.
6. Click Uninstall Application to mark the application as obsolete.

Figure 4.3 The Uninstall Application dialog box

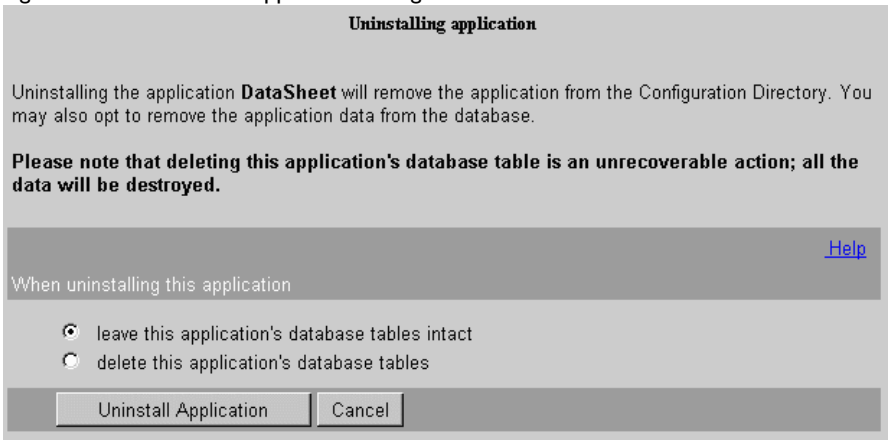

Now you and your end users cannot access the process instance detail for the application. End users cannot initiate new process instances or continue working on existing work items.

Viewing the Application Logs

PAE applications use several log files, with each Netscape Application Server in a cluster having its own set of log files. PAE administrators can view these application logs that are maintained for each application on a per-server basis:

- warning log (`warningLog.html`) -- shows warnings that affect the application.
- error log (`errorLog.html`) -- shows the errors that have occurred with the application.

These files are located at *`serverRoot`*`/builder/support/log`.

To view an application's logs, follow these steps:

1. Go to Process Administrator's home page at `http://`*`yourServer`*`/Administrator.apm`.
2. Click the Applications tab. This displays the list of applications in the cluster.
3. For the application whose logs you want to view, choose the View Logs command from the drop-down list for that application.
4. Click Apply.
5. Click View Log File.
6. If a dialog box prompts you to enter a user ID and password, enter the ones that you defined for your Administration Server. For example, `admin/admin`. When they are authenticated, the View Application Logs page is displayed.
7. From the View Application Logs page, choose the log you want to view from the drop-down list.
8. Indicate how many log records you want to view. The log file display starts at the end of the file, so indicating "12" would show the last 12 records in the file.
9. Click View Log to see the log file.

Exporting and Deleting Data

As PAE administrator, you can export or delete application user data that's no longer in use. You can only export and delete data from closed process instances. You cannot delete or export work items separately from the process instances that they are part of.

You can choose to export only, export and then delete, or delete only. You can perform the action for a single application or for all applications at once.

You can access this function through the Export and Delete Data button at the bottom of the main Applications page.

To export or delete user data, follow these steps:

1. Go to Process Administrator's home page at `http://yourServer/Administrator.apm`.
2. Click the Applications tab.
3. Click the Export and Delete Data button at the bottom of the page below the list of applications.

   The window pictured in Figure 4.4 appears:

Figure 4.4 The Export and Delete Data page

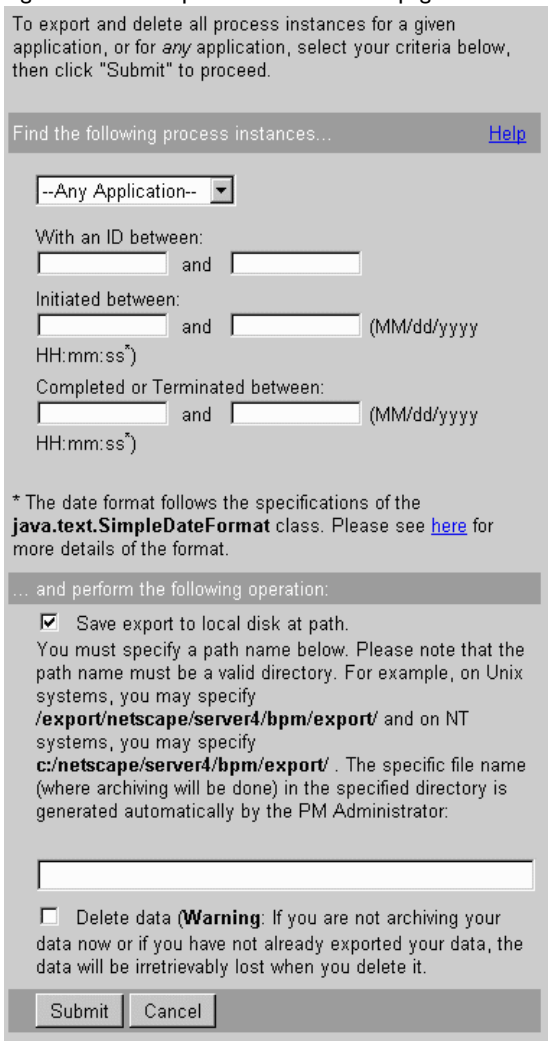

4. Select a specific application or “Any Applications” from the drop-down list.

5. Enter any search criteria.

   You search for process instances based on one or more of the following criteria:

   - process instance IDs
   - time the process instance was created
   - time the process instance was completed or terminated

6. Check "Save export to local disk at path" to export the data.

NT On Windows NT, the disk must be local. You cannot use an NFS-mounted drive.

7. Enter the absolute physical path to the folder in your local file system that you want to use for archiving. If you checked the "Save export..." checkbox, you must enter a path or you get an error when you try to perform the operation.

8. Check "Delete data" to delete the data. You will lose any data that you haven't exported previously or as part of this operation.

   You should export all data before deleting it.

9. Click Submit to export (and also delete if you checked the delete box) all the application data that meets your criteria.

Note You can delete data without archiving it, although this is a very risky action to take. This action is not recoverable.

Exported process instance data is written to flat files in a directory the administrator specifies. Each process instance's data is exported to a separate file. The file name for each process instance is generated automatically in the following format:

```
process_instance_applicationName_processInstanceID_year_month_day_hour_
minute_second
```

For example:

```
c:\tmp\process_instance_CreditHistory_10_1999_10_11_10_31_51.xml
```

Here, the archival file was written to the directory `c:\tmp\`, the application name is CreditHistory, the process instance ID is 10, and the archival was done on October 11, 1999 at 10:31 and 51 seconds.

Exported data is stored in an xml file with the following format:

```
<PROCESS_INSTANCE>

<INSTANCE_DATA>

This section contains data from the process_instance table.

</INSTANCE_DATA>

<HISTORY_EVENT_LIST>

<HISTORY_EVENT>

This section contains data from one history event in the history table.
There may be more than one history event listed here.

</HISTORY_EVENT>

</HISTORY_EVENT_LIST>

<DATA_ELEMENT_LIST>

This section contains a list of all data elements and their
corresponding values. This section also contains custom field data.

</DATA_ELEMENT_LIST>

</PROCESS_INSTANCE>
```

Chapter

5

Databases

Netscape Application Server: Process Automation Edition (PAE) stores user data in a relational database. The database provides a centrally available repository for user data for each application in a cluster, plus a common set of tables that all applications can share. In addition, there are adapters for each type of database that handle mapping between PAE types and the data types used by each kind of database, so process designers can design process definitions that are database independent.

This chapter describes these topics:

- Overview
- Using Databases With PAE
- Setting Up Your Database

Note This chapter only provides general guidelines for using databases with PAE. For specific situations and error conditions, consult your database administrator (DBA). For example, if you encounter a Sybase log full error, you should refer the problem to your database administrator for a solution.

Overview

Netscape Process Manager uses an Oracle or Sybase database to store process instance data, so you must have a database available to you on your local system or available over a network on another system. If you are using a remote database, you must have the appropriate database client installed on your local system. You need to have enough privileges to create tables, write to tables, and to create views.

Process Manager supports the following databases:

- Oracle 8.0.5
- Sybase 11.9.2 (server), 11.1.1 (client)

There are several ways a PAE end user and administrator interact with database tables. Sometimes they need read-only access, but sometimes they add new data or update existing data. Note that the database access is transparent to the end user or administrator. They are only aware that they are viewing, adding, or modifying process instance data.

End users are accessing a database when they perform these operations:

- Create a new process instance: Adds new data to all tables in the database.
- Work on a work item: Adds new application-specific data and updates some of the cross-application tables.
- View a process instance detail: Reads from several tables.
- View a process instance's history: Reads from the history table.

Administrators are accessing a database when they perform these operations:

- View a process instance detail: Reads from several tables.
- View a process instance's history: Reads from the history table.
- Request process instance statistics: Reads from several tables.
- Archive data: Updates some tables by flagging the database records that match the archiving criteria.

- Delete data through the Archive and Delete page: Updates some tables by removing the database records that match the deletion criteria. This is not recoverable.
- Delete data when uninstalling an application: Deletes the application-specific table for the application. This is not recoverable.
- Delete data when removing a cluster: Deletes the application-specific and the cross-application tables. This is not recoverable.

At each database operation, PAE collects data into implicit database transactions to protect application data in the database. Transaction-based processing allows rollbacks and automatic database cleanup if any PAE operation fails.

Note Each of these transactions generates an entry in the database transaction log, so make sure you have plenty of space available for the log. If you are using a Sybase database, it disconnects automatically when there is not enough space for additional transaction log entries. You can increase the transaction log size for your Sybase database to help avoid this problem.

Using Databases With PAE

To provide consistency across all applications in a cluster, all applications share a single relational database.

This section includes these topics:

- Cross-Application Tables
- Primary & Foreign Keys
- Database Views
- Database Users
- Database Adapters
- Storing Digital Signatures
- Setting Up Your Database
- Database Connections

Cross-Application Tables

The database has five common, cross-application tables that are shared across all applications and one application-specific table for each application (this value is used in the DB Application Table field in the application's definition in Process Builder). The cross-application tables contain standard data for each process instance and work item, while the application-specific tables contain the user data that is unique to each application.

The cross-application tables are:

- process_instance table: Maintains process_instance data such as each process instance's state, application, and creation and modification dates. It also tracks who initiated the process instance and who it last worked on it.
- work_item table: Tracks work item data such as each work item's expiration date and who it is currently assigned to.
- history table: Maintains all historical data for each process instance.
- wf_id_range table: Generates a new unique process instance ID whenever an end user initiates a new process instance. This table has a single field, `next range`, indicating the next ID to create.
- wf-blobs table: Stores digital signatures.

Table 5.1 The database tables

process_instance	work_item	history	wf_blobs
wf_instance_id	wf_instance_id	wf_instance_id	wf_instance_id
wf_creation_date	wf_fork_id	wf_fork_id	wf_field_cn
wf_creator_dn	wf_node_cn	wf_event_id	wf_blob_content
wf_title	wf_wi_state	wf_event_type	
wf_priority	wf_expired_flag	wf_event_date	
wf_entry_cn	wf_wi_cdate	wf_wi_cdate	
wf_exit_cn	wf_exp_date	wf_wi_exp_date	
wf_last_modified	wf_user_dn	wf_user_dn	
wf_pi_state	wf_target_cn	wf_app_cn	

Table 5.1 The database tables

process_instance	work_item	history	wf_blobs
wf_app_cn	wf_performer_url	wf_comment	
wf_observer_url			
wf_performer_url			
wf_archived_flag			

Note In a Sybase database, the maximum number of bytes per row is 1962. If you create tables with varchar, nvarchar, or varbinary columns whose total defined width is greater than 1962 bytes, a warning message appears, but the table is created.

Primary & Foreign Keys

The database tables are set up to share the process instance, `wf_instance_id`, as a key. The process_instance table uses the `wf_instance_id` field as the primary key. The `wf_instance_id` field in other tables is considered as a foreign key. The work_item table uses a concatenated key of `wf_instance_id`, `wf_fork_id`, and `wf_node_cn` to precisely identify a work item in a parallel processing application, where there may be several work items for a specific process instance.

Database Views

To allow simplified queries, PAE provides some standard views into the database information. These offer a predefined subset of the data available in the database.

PAE constructs these views:

- closed_process: Uses the process_instance table in this query:

```
SELECT wf_instance_id, wf_creation_date, wf_creator_dn, wf_title,
wf_priority, wf_entry_cn, wf_exit_cn wf_node_cn, wf_last_modified,
wf_pi_state, wf_app_cn, wf_observer_url, wf_performer_url,
wf_archived_flag
FROM process_instance WHERE wf_pi_state >= 4
```

- worklist: Uses the process_instance and work_item tables in this query:

```
SELECT  wi.wf_instance_id, wi.wf_performer_url,
wf_observer_url, wf_fork_id,
wf_node_cn, wf_wi_state,
wf_expired_flag, wf_wi_cdate,
wf_exp_date, wf_user_dn, wf_title,
wf_priority, wf_creation_date,
wf_creator_dn, wf_entry_cn,
wf_last_modified, wf_pi_state,
wf_target_cn, wf_app_cn
FROM process_instance pi, work_item wi
WHERE wi.wf_instance_id = pi.wf_instance_id
```

- application-specific view: Uses the process_instance and the application-specific tables in a query such as this:

```
SELECT  pi.wf_instance_id, fieldrole, datefield, textfield, radio,
wf_creation_date, wf_creator_dn,
wf_last_modified, wf_app_cn,
wf_ pi_state, wf_entry_cn, wf_title,
wf_priority, wf_exit_cn, wf_archived_flag,
wf_performer_url,wf_observer_url
FROM user1.process_instance pi, child app
WHERE pi.wf_instance_id = app.wf_instance_id
```

- searchlist view: Uses the process_instance table in a query such as this:

```
SELECT pi.wf_instance_id wf_instance_id,
wf_creation_date,  wf_creator_dn, wf_title, wf_priority,
wf_entry_cn, wf_exit_cn, wf_last_modified, wf_pi_state,
wf_app_cn, wf_observer_url, wf_archived_flag, wf_fork_id,
wf_node_cn, wf_wi_state, wf_expired_flag, wf_wi_cdate,
wf_exp_date, wf_user_dn, wf_target_cn, wi.wf_performer_url,
wf_deferred_flag from process_instance pi, work_item wi
where pi.wf_instance_id = wi.wf_instance_id (+))
```

Database Users

There are three kinds of database users:

- Administrator

  Has create and write access to all cross-application and application-specific tables and views. This is the database user that you identify when creating a cluster. This administrator is responsible for setting up your database.

- Business Manager

  Also has create and write access to all cross-application and application-specific tables and views. This business manager is responsible for process instance-specific database issues. See the *Business Manager's Guide* for more information.

- application-specific user

  Has create and write access to their own application's tables and views. Has read-only access to cross-application tables and views. Has no access to other application-specific tables.

  If you have an application-specific database user, the Process Builder designer identifies this as the DB User in the application's properties window when deploying the application.

Database Adapters

Each type of database supported by PAE has its own adapter. After the database type is determined, the adapter handles storing data in the database and takes care of such database-specific characteristics as type mapping.

Because type mapping in PAE is handled by the database adapters, Process Builder designers can use a single set of types that can be correctly interpreted for each type of database. Designers do not need to concern themselves about how their data types are stored in a given database.

Note Designers who write their own customized scripts cannot use certain transaction control statements in their scripts such as begin, commit, and abort. In addition, if they are using an Oracle database, they cannot use DDL in their script statements. Also, Process Builder does not check for all reserved words for every database. If a custom script contains a reserved word, your application deployment may fail.

Storing Digital Signatures

PAE has a special database table for storing digital signatures, wf_blobs.

Table 5.2 Fields in the wf_blobs database table

wf_blobs
wf_instance_id
wf_field_cn
wf_blob_content

You can perform `dbadmin` queries against this table to verify digital signatures as required. For example, you could execute this query

```
SELECT * from wf_blobs WHERE wf_instance_id = <ID> and wf-field_cn = '<sig-field>'
```

to obtain the user ID, user name, and digital signature for a specific process instance. Note that to do this, you will need the process instance ID.

Setting Up Your Database

PAE requires the use of an Oracle or Sybase database. You can use whatever database you already have installed, or you can install a new database for your PAE data. You, or the database user you identify, must have create and write permissions for tables and views.

Note that if you are using Sybase, a database already exists for PAE's use. The database administrator of your enterprise should have created this database before creating a cluster. The database administrator must make sure the PAE administrative user has create table/create view privileges on that database. Sybase databases are no longer created during cluster creation.

Information you need to know before you attempt to create a new cluster or to join an existing cluster are:

- Database server type (Oracle or Sybase)
- Database server identifier

- Valid database user ID (the administrator's user ID). Note that this user cannot be the same as the database owner for Sybase databases.
- Password for that user
- Database (for Sybase only)

See the *Installation Guide* and the *Netscape Application Server Administration Guide* for more information about setting up your database. You can locate Netscape Application Server documentation in the following location:
serverRoot/nas/docs/index.htm

Database Connections

Database connections are maintained through Netscape Application Server. You can set connection parameters in Netscape Application Server (NAS) Administrator.

To set the connection parameters, perform the following steps:

1. From Netscape Console or the Windows NT Start menu, launch NAS Administrator.
2. On the NAS Administrator toolbar, click the Database button to open the Database window.
3. In the left pane of the Database window, click the database for which you want to adjust the timeout parameter.

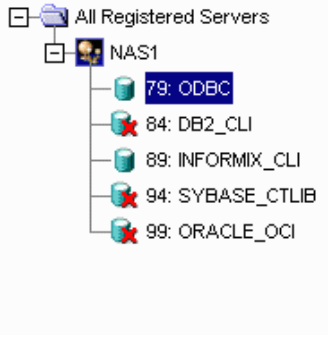

4. In the right pane of the Database window, in the Connection Timeout field, enter the number of seconds.

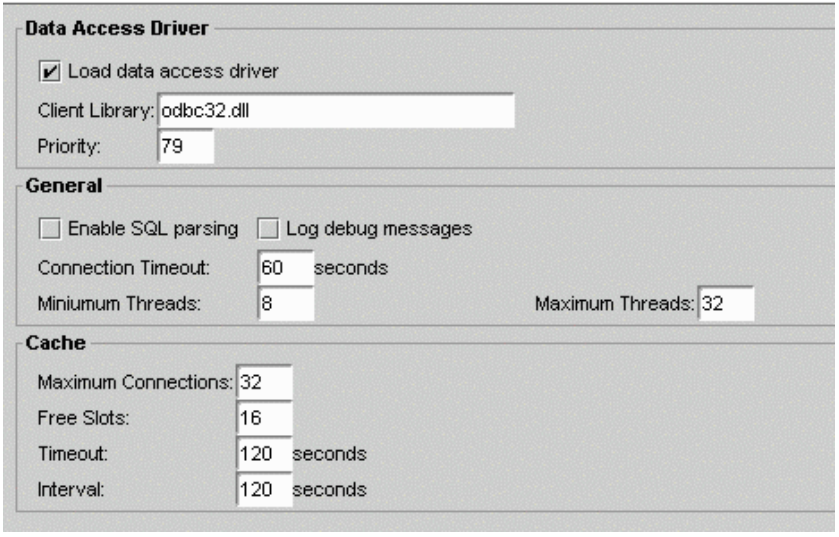

You can also set the minimum and maximum number of threads available for database connections. The thread parameters determine how many threads NAS allocates for asynchronous database queries. Such threads are usually used for queries returning a large number of rows and allowing the application to do other tasks while waiting for the query to finish. Asynchronous database queries are not supported by JDBC 2.0, a Java programming interface used to build on top on database drivers.

The default thread allocations are adequate for most applications. If an application developer uses many asynchronous queries, you might want to increase the maximum number of available threads. Keep in mind that each thread does use a small stack allocation and pulls from the total number of available system threads. Therefore, if an application does not use any asynchronous queries, you can increase performance by setting the maximum available threads to zero.

To set the thread parameters, perform the following steps:

5. In the right pane of the Database window, in the Minimum Threads field, enter the number of threads.

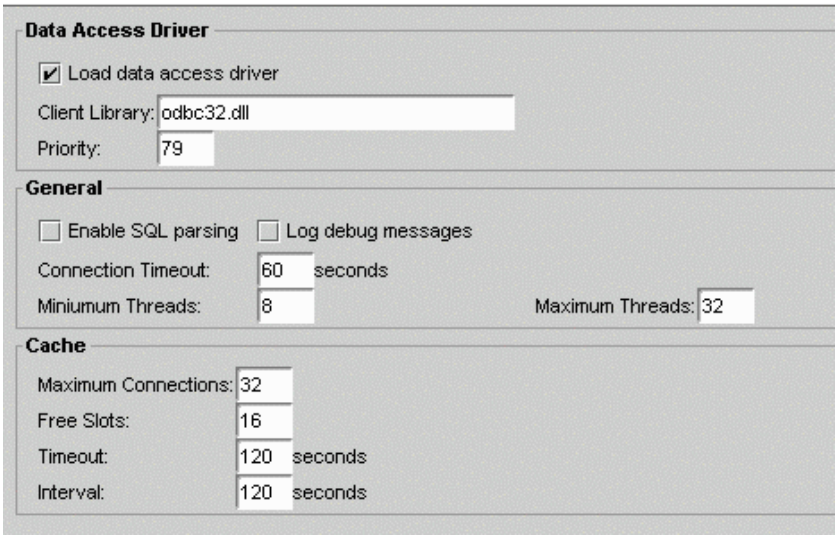

6. In the right pane of the Database window, in the Maximum Threads field, enter the number of threads.

The database cache is an array used to hold active and recently used database connections. NAS adds database connections to cache when an application creates a database connection.

While the application is using that database connection, NAS marks that connection "in use." Once the database operations are finished, the server marks the database connection "free." The cache then holds the free connection in the cache for a configured period of time. This allows the server to use the free cached connection and quickly handle a new request to the same database. Once a free connection exceeds the timeout, a cleaning thread removes the connection from the cache and opens a slot for a new connection to be cached.

You can adjust the following cache parameters:

- the maximum number of connections allowed in the cache
- the number of slots held solely for free connections

- the timeout limit, in seconds, for free connections
- the interval, in seconds, at which the cache cleaner thread removes timed-out free connections

The default values are adequate for most applications, so adjustments are not usually required for initial application installations.

NAS dynamically adjusts the cache up to the maximum number of allowable connections. If there are no connections to cache, the array is allocated to zero spaces.

7. In the right pane of the Database window, under Cache, enter values for the following parameters:

- maximum connections
- free slots
- timeout
- interval

8. Click Apply Settings to save the changes to NAS.

It is perhaps worth noting here that Process Express, Process Administrator, and Process Business Manager also each use a database connection. Whenever you are tracking the number of connections in use at a given time, remember that two connections are always taken by these components.

Glossary

activity A step in an application where an assignee needs to perform an action. This is referred to as a *work item* in Process Express.

application In Process Builder, the application the designer creates to handle a PAE process.

assignee The person assigned to an activity for a particular process instance.

automated activity A step in an application where an action is performed automatically, without an assignee.

builder The person who creates the application using Process Builder.

child process In subprocesses, the subordinate process that is called by the main or parent process.

class ID The identifier of a group of fields with certain common properties.

CGI Common Gateway Interface. The specification for communication between an HTTP server and gateway programs on the server. Allows web interfaces to databases and enables the dynamic generation of HTML documents by gateway programs.

cluster The association of a configuration directory, a corporate user directory, one or more application servers, and a database.

configuration directory The Netscape Directory Server where PAE cluster and application information is stored.

content store The Enterprise Server folder in which file attachments and the user and password needed to access them are stored.

corporate user directory The LDAP directory service used to store the user and group information for a corporation. Process Builder reads from it when a designer sets up an application's users and groups and assigns work items. Process Express, Process Administrator, and Process Business Manager read from it when identifying a work item's assignee or when an assignee wants to delegate a task to another user.

creator The person who initiates a process instance. In Process Express, Process Administrator, and Process Business Manager, called the *initiator*.

database The relational database that stores the information generated by process instances. For example, the database could be Informix, Oracle, or Sybase.

decision point A point at which a process map branches depending upon conditions defined in the decision point.

deploy To copy an application stored locally to a cluster. It can be deployed for storage only, or it can be deployed for testing or production. If it is deployed for testing or production, the application information is deployed to the configuration directory, and the application is activated on the application server.

entry point A point in the process where a user can initiate a process instance.

exception handler Used in subprocesses, a step in an application that allows the administrator to intervene manually if errors occur in the interaction between a parent and child process.

exit point A point in the process where the process ends.

extranet An extension of a company's intranet onto the Internet, to allow customers, suppliers, and remote workers access to the data.

form A part of an application a user fills out to complete a process instance, or uses to view information on a process.

group A set of users defined in the corporate user directory or within an individual application to whom the designer can assign an activity, or work item,

HTML HyperText Markup Language. A markup language (derived from SGML) used to create web documents.

HTTP HyperText Transfer Protocol. A protocol for communication between web clients and servers.

initiator The person who initiates a process instance. In Process Builder, called the *creator*.

intranet A network which provides similar services within an organization to those provided by the Internet outside it but which is not necessarily connected to the Internet.

nested parallel process A parallel process nested within a larger parallel process. The activities in the nested process are considered to be part of the nested process and not the larger process.

parallel processing A step in an application that branches between two or more branches so that two or more activities can execute in parallel.

parent process In subprocesses, the main process that calls the subordinate or child process.

participant A user of Process Express.

Process Administrator The component of PAE that administrators use to administer PAE and PAE applications. Process Administrator is the IT administrator's interface described in the *Administrator's Guide*. Process Business Manager is the business manager's interface, described in the *Business Manager's Guide*.

Process Express The component of PAE that end users use to initiate process instances, complete work items, and search for process instances.

Process Builder The component of PAE that designers use to build and deploy applications.

Process Business Manager The component of PAE that business managers use to administer PAE work items and process instances. Process Business Manager is the business manager's interface, described in the *Business Manager's Guide*. Process Administrator is the IT administrator's interface described in the *Administrator's Guide*.

Process Engine The PAE software component that contains Process Express, Process Administrator, Process Business Manager and the internal software that runs PAE.

process A process is a series of activities, or *work items*, that can be completed by end users using Process Express.

process instance A particular example of a PAE process; for example, in a time off process, a process instance would be a particular request by an employee for vacation time off for a specific period of time.

PAE Netscape Application Server: Process Automation Edition, a process management solution.

process map The visual representation of the process a PAE application handles.

processing branch A set of activities that progress from a given split to its corresponding join. Also called a *thread*.

property	An attribute of an item or component used in an application that contains information about the item. For example, an activity has properties containing information such as the name of the activity and what script is run when it is completed.
role	A role is the part a user plays in a specific process instance.
script	In PAE, a script is a JavaScript script that can, but does not have to, include a function.
Simple Workflow Access Protocol (SWAP)	A protocol that allows process management systems to instigate, monitor, and control workflow processes on other process management systems.
subprocess	A fully functional process that is called from within another process. The process that calls the subprocess is the parent process and the subprocess is its child process.
trusted user	A group in a subprocess. By adding a user to this group that matches the AppUserID in the parent process, you can set up a reliable secure handshake between two specific applications.
transition	The links between steps in a process. On the process map, they are represented by lines with arrows that lead from one item to another.
work item	A work item is an individual task in a process as it appears to the end user on a work list. This is referred to as an **activity** in Process Builder.

Index

A

activities 11
Admin Data Auth 26
Admin IT Auth 26
Administration Server log files 29
application stages 52
application status 52
applications 51
  archiving data 58
  autostart 55
  closing 57
  defined 52
  deleting data 58
  deploying 54
  forms 19
  initializing 55
  log files 28, 60
  managing 54
  reopening 57, 58
  starting 56
  stopping 56
  uninstalling 53, 58
  viewing logs 60
application-specific database tables 68
  deleting 58
archiving
  application data 58
  process instances 61
asynchronous database queries 74
autostarting applications 55

B

Base DN 20, 36, 45
Basic Auth 26
Bind DN 20, 36, 37, 42, 45
Bind Password 36

C

cache 75
  parameters, adjusting 75
certificates 18
Change Cluster page 44
changing the corporate user directory 43
Closed stage 53
Cluster Description 35
Cluster Display Name 35
Cluster DN 35
cluster management
  forms 19
  Process Administrator logs 46
Cluster Management page 42
clusters 10, 22, 31
  accessing Directory Server information 45
  changing information 43
  configuration directory information for 37, 41
  corporate user directory information for 36
  creating 33
  database information for 38
  defined 32
  deleting 47
  general information 35, 41
  joining 40
  Mail Server information for 39
  managing 42
cn (common name) 21
configuration directory 10, 12
Configuration Directory Port Number 37
configuration files

Enterprise Server 26
Process Builder 27
conventions used in this book 7
corporate user directory 10, 12
changing 43
Create Cluster page 33
Create or Join Cluster page 33, 40
creating a cluster 33
customized scripts, using with databases 71

D

database connection parameters
setting 73
database connections
caching 75
threads 74
Database Server Identifier 38
Database Server Type 38
database tables
application-specific 68
cross-application 68
fields 68, 72
history 68
process_instance 68
wf_blobs 68, 72
wf_id_range 68
work_item 68
databases 10
adapters 71
cross-application tables 68
Database Server Identifier 38
foreign keys 69
primary key 69
users 38, 70
using customized scripts with 71
views 69
DB Application Table 68
DB user 71
dbadmin 72
deleting a cluster 47
deleting database tables 58
Deploy Applications dialog box 53
deploying an application 54
Design stage 53
digital signatures 18, 72
directories 12
Directory Manager 20
Directory Server
accessing information from a cluster 45
Directory Server Host Name 36, 37, 41, 45
Directory Server log files 29
Directory Server Port Number 36
Directory Server terms 19
Directory Suffix 20
Distinguished Name 20
DN. See Distinguished Name

E

Enterprise Server
SSL-enabled 18
Enterprise Server log files 29
error log 60

F

fonts used in this book 7
foreign keys 69

H

history database table 68
HTML versions of guides 8

I

initializing an application 55

J

Join Cluster page 40

joining a cluster 40

L

LDAP terms 19
LDAP URLs 22, 27
LDAP. See Directory Server
log files
  applications 28, 60
logs 28
  Administration Server 29
  Directory Server 29
  Enterprise Server 29
  Process Administrator 28, 46
  Process Builder 28

N

Netscape Application Servers 17

O

o (organization) 21
obj.conf file 26
Obsolete applications 58
Obsolete stage 53
Open stage 53
Oracle 66
ou (organizational unit) 21

P

PAE components 10
PDF versions of guides 8
preferences.ini file 22, 27, 34, 43, 47
primary key 69
Process Administrator logs 28
Process Builder
  DB User 71
  Deploy Applications dialog box 53
  logs 28
  preferences.ini file 27
process instances 10
  archiving 61
  deleting 61
Process Manager directories 12
process_instance database table 68

R

Removing Timer Agents 47
root DN 20
root entry 20

S

search base 20
security 18
Server Identifier
  database 38
SMTP Port 39
SMTP Reply To 39
SMTP Server 39
SSL-enabled Enterprise Servers 18
stages, applications 52
Started status 53
status, applications 52
Stopped status 53
Stored Only stage. See Design stage
styles, in this book 7
Sybase 66
Sybase database table limits 69

T

terms, in this book 7
Testing stage 53
thread parameters
  setting 74
threads

database connections 74
type mapping 71

U

uid (user ID) 21
Uninstall Application dialog box 59
uninstalling a closed application 53
unrestricted user 20
URLs
format in manual 7

V

View Process Administrator Logs page 46

W

warning log 60
wf_blobs database table 68, 72
wf_id_range database table 68
work items 11
work_item database table 68