

IBM FileNet P8
Version 4.5.1

*Plan and Prepare Your Environment for
IBM FileNet P8
for installation on Solaris with Oracle,
Oracle WebLogic Server, and Sun Java
System Directory Server*



IBM FileNet P8
Version 4.5.1

*Plan and Prepare Your Environment for
IBM FileNet P8
for installation on Solaris with Oracle,
Oracle WebLogic Server, and Sun Java
System Directory Server*



Note

Before using this information and the product it supports, read the information in "Notices" on page 77.

This edition applies to version 4.5.1 of IBM FileNet Content Manager (product number 5724-R81), version 4.5.1 of IBM FileNet Business Process Manager (product number 5724-R76), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2001, 2010.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

ibm.com and related resources. v

How to send your comments. v

Contacting IBM vi

Planning and preparing for IBM FileNet

P8 installation 1

Planning the installation 1

FileNet P8 Platform sample architecture 1

Baseline configuration with expansion components. 2

Sample development architecture 2

Sample demonstration architecture 3

Planning your deployment 3

Planning Content Engine deployment 3

Application server planning considerations 4

WebSphere Application Server and

WebLogic Server - Content Engine

multi-server deployment 5

Definition of installation roles 5

Using the installation and upgrade worksheet 7

Running the Customize Worksheet macro 8

Autofiltering and sorting the Worksheet 8

Performing the required installation preparation tasks 9

IT administrator installation tasks 10

Creating Content Engine operating system accounts 10

Creating Process Engine operating system accounts 13

Configuring UNIX 15

Configuring UNIX for IBM FileNet P8 servers (all components) 16

Configuring Content Engine servers (all UNIX) 16

Configuring Content Search Engine servers (all UNIX) 17

Configuring Process Engine servers (all UNIX) 17

Configuring Process Engine servers (Solaris) 19

Configuring Application Engine (Solaris) 23

Configuring the network 23

Prerequisites to configuring your network 23

Synchronizing time and date 24

Preparing storage areas for object stores 24

Configuring file servers for file storage areas 25

Users and groups 26

Configuring the remote access protocol on the client machine 28

Preparing storage areas to support content search functionality. 28

Security administrator installation tasks 29

Security planning considerations 29

Configuring directory server. 31

Configuring Sun Java System Directory

Server 31

Specifying IBM FileNet P8 accounts 33

Creating Content Engine directory server accounts 33

Creating Process Engine directory server accounts 36

Creating Application Engine or Workplace XT accounts 37

Creating Content Search Engine accounts 38

Database administrator installation tasks 39

Creating Content Engine database accounts. 40

Creating Process Engine database accounts 41

Preparing Oracle server 42

Oracle planning considerations 42

Verifying that Oracle server is installed for IBM FileNet P8 43

Running Process Engine pre-installation scripts (Oracle) 50

Configuring Oracle client for Process Engine 55

Installing the Oracle client 55

Verifying the ability to connect to the database 56

Verifying the Process Engine database connection (Oracle). 56

Application Server administrator installation tasks. 57

Creating Content Engine application server accounts 58

Configuring WebLogic Server for Content Engine 58

Configuring WebLogic Server for Application Engine or Workplace XT 60

Configuring Process Engine clients for ORB 61

Configuring the documentation server 61

Appendix A. Preparing non-English environments for installing IBM FileNet P8 Platform. 63

Application Server administrator 63

Configuring character encoding (WebLogic Server) 63

WebLogic Server. 64

Security administrator 64

Extended characters and user names 64

Database administrator 64

Installing Oracle server 64

Installing Oracle client. 65

NLS_LANG environment variable - charset 65

UNIX 66

Setting the NLS_LANG environment

variables 66

IT administrator 67

Operating system considerations 67

UNIX 68

| | |
|---|----|
| Configuring support for other languages in X Windows UNIX system | 68 |
| Setting the LANG, LC_TIME, and LC_MESSAGES environment variables | 68 |
| IBM FileNet P8 administrator | 69 |
| Process Engine services | 69 |
| Configuring Process Task Manager for Application Engine and Workplace XT | 69 |
| Limitations on installing in a non-English environment | 69 |

| | |
|---|-----------|
| Appendix B. IBM FileNet P8 ports | 71 |
|---|-----------|

| | |
|--------------------------|-----------|
| Notices | 77 |
| Trademarks | 79 |

| | |
|------------------------|-----------|
| Index | 81 |
|------------------------|-----------|

ibm.com and related resources

Product support and documentation are available from [ibm.com](http://www.ibm.com).

Support and assistance

Product support is available on the Web. Click Support from the product Web site at:

FileNet Content Manager Support

<http://www.ibm.com/software/data/content-management/filenet-content-manager/support.html>

Information center

You can view the product documentation in an Eclipse-based information center that you can install when you install the product. By default, the information center runs in a Web server mode that other Web browsers can access. You can also run it locally on your workstation. See the information center at <http://publib.boulder.ibm.com/infocenter/p8docs/v4r5m1/index.jsp>.

PDF publications

You can view the PDF files online using the Adobe Acrobat Reader for your operating system. If you do not have the Acrobat Reader installed, you can download it from the Adobe Web site at <http://www.adobe.com>.

See the following PDF publications Web sites:

| Product | Web site |
|---|---|
| Product Documentation for FileNet P8 Platform | http://www.ibm.com/support/docview.wss?rs=3278&uid=swg27010422 |

"How to send your comments"

Your feedback is important in helping to provide the most accurate and highest quality information.

"Contacting IBM" on page vi

To contact IBM customer service in the United States or Canada, call 1-800-IBM-SERV (1-800-426-7378).

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Send your comments by using the online reader comment form at https://www14.software.ibm.com/webapp/iwm/web/signup.do?lang=en_US&source=swg-rcf.

Consumability survey

You are invited to tell IBM how to improve the consumability of software products. If you want to help IBM make IBM® FileNet® P8 easier to use, take the Consumability Survey at <http://www.ibm.com/software/data/info/consumability-survey/>.

Contacting IBM

To contact IBM customer service in the United States or Canada, call 1-800-IBM-SERV (1-800-426-7378).

To learn about available service options, call one of the following numbers:

- In the United States: 1-888-426-4343
- In Canada: 1-800-465-9600

For more information about how to contact IBM, see the Contact IBM Web site at <http://www.ibm.com/contact/us/>.

Planning and preparing for IBM FileNet P8 installation

To prepare to install IBM FileNet P8 Platform, you must review the planning information before you begin. You must also complete the prerequisite tasks assigned to the various roles.

“Planning the installation”

You must review the installation planning information before your IBM FileNet P8 Platform installation so that you know what kind of deployments are supported, understand how the tasks in the installation tasks are organized by role, and know how to use the Installation and Upgrade Worksheet.

“Performing the required installation preparation tasks” on page 9

To efficiently carry out the required installation preparation tasks, you must assign your staff to carry out the tasks that are organized by administrative role.

Planning the installation

You must review the installation planning information before your IBM FileNet P8 Platform installation so that you know what kind of deployments are supported, understand how the tasks in the installation tasks are organized by role, and know how to use the Installation and Upgrade Worksheet.

“FileNet P8 Platform sample architecture”

You can distribute IBM FileNet P8 Platform components and expansion products across a variety of machines.

“Planning your deployment” on page 3

Content Engine, Process Engine, and Application Engine, or Workplace XT, can all be installed as standalone, single instance applications, or as multiple instances with different names on the same server, or in a cluster.

“Definition of installation roles” on page 5

The tasks in this guide and the rows in the Installation and Upgrade Worksheet are organized by administrative roles. Your organization might have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.

“Using the installation and upgrade worksheet” on page 7

The Installation and Upgrade Worksheet is a Microsoft[®] Excel spreadsheet (p8_worksheet.xls). The worksheet describes the properties and parameters required to complete IBM FileNet P8 installation, upgrade, and configuration programs, and provides a way to record the values you assign to these properties and parameters.

FileNet P8 Platform sample architecture

You can distribute IBM FileNet P8 Platform components and expansion products across a variety of machines.

“Baseline configuration with expansion components” on page 2

A production environment includes required components plus one or more optional expansion components.

“Sample development architecture” on page 2

To provide an application development configuration, you can install the components that the application is based on and also the required API toolkits on developer workstations.

“Sample demonstration architecture” on page 3

You can install required and any expansion components on a single server in order to support demonstrations, proof-of-concepts, and single-server development. Because Enterprise Manager runs only on Windows, a single-server installation that includes this administration tool must be on a Windows server.

Baseline configuration with expansion components

A production environment includes required components plus one or more optional expansion components.

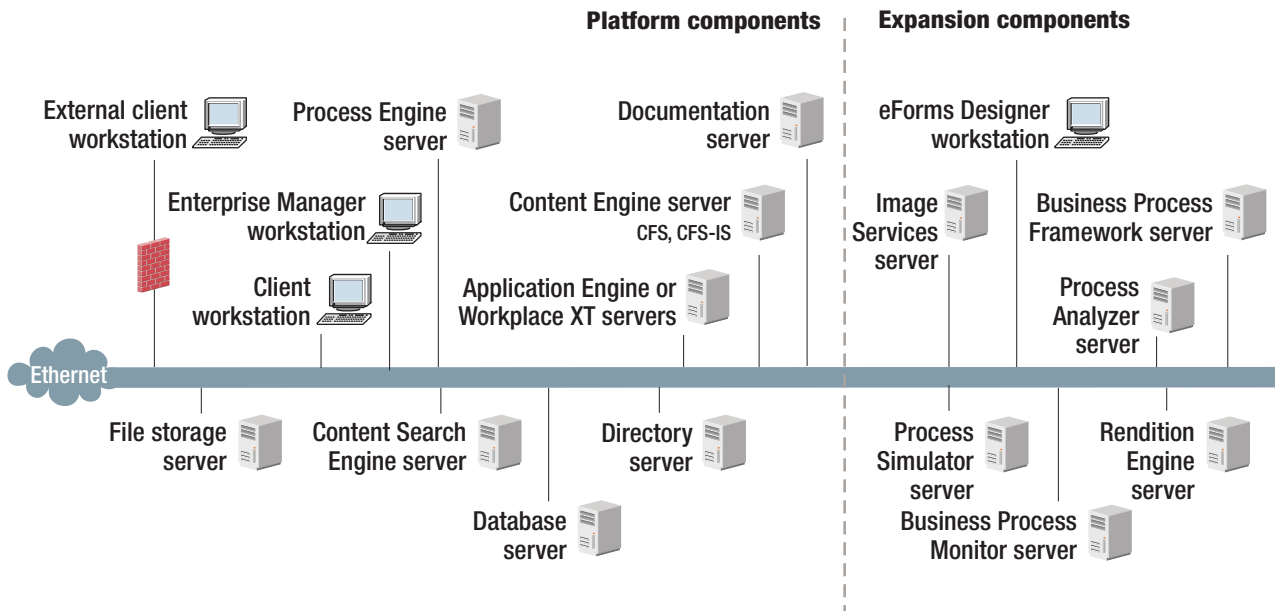


Figure 1. Baseline configuration with expansion components

To understand this graphic, keep in mind the following details:

- You must set up Content Engine, Application Engine, and the documentation server on application servers.
- IBM FileNet P8 Platform supports several different operating systems. However, Enterprise Manager runs only on Windows.
- The graphic does not show components in a high availability or clustered configuration.
- For information about collocating components on a single server, see *IBM FileNet P8 Hardware and Software Requirements*.
- Not all expansion components are shown, for example IBM FileNet P8 Portlets and Image Services Resource Adapter. Check with your service representative for availability of other expansion products.

Sample development architecture

To provide an application development configuration, you can install the components that the application is based on and also the required API toolkits on developer workstations.

Refer to the IBM FileNet P8 Developer Help topic **IBM FileNet P8 Documentation** → **Developer Help** → **Developer Roadmap** → **Introduction** for information on setting up your development environment and installing the IBM FileNet P8 API toolkits.

Sample demonstration architecture

You can install required and any expansion components on a single server in order to support demonstrations, proof-of-concepts, and single-server development. Because Enterprise Manager runs only on Windows, a single-server installation that includes this administration tool must be on a Windows server.

Installing many FileNet P8 components on the same server is likely to affect system performance. For more information, see the *IBM FileNet P8 Hardware and Software Requirements* and the *IBM FileNet P8 Performance Tuning Guide*.

Planning your deployment

Content Engine, Process Engine, and Application Engine, or Workplace XT, can all be installed as standalone, single instance applications, or as multiple instances with different names on the same server, or in a cluster.

“Planning Content Engine deployment”

You can deploy Content Engine in either a stand-alone, multi-instance single-server, or cluster environment.

“Application server planning considerations” on page 4

To prepare your application servers, you must carry out certain tasks including creating administrative accounts, specifying environmental variables, and making sure your JDBC drivers are at a supported level.

Planning Content Engine deployment

You can deploy Content Engine in either a stand-alone, multi-instance single-server, or cluster environment.

Standalone deployment

When you deploy Content Engine as a stand-alone application, you configure a single application server. You must configure your Content Engine instances and deploy those Content Engine instances on a single server, using a single directory for the configuration files.

Multi-instance single-server deployment

When you deploy multiple Content Engine instances on a single server, you configure a single application server.

Each Content Engine instance is isolated from the others, and there is no exchange of information between the instances. For example, you can dedicate a Content Engine instance for use by the Research and Development department, and you can dedicate a Content Engine instance for use by the Human Resources department. Each instance has its own set of object stores.

You must perform all the procedures for each instance on a single server, using a different directory for the configuration files for each instance. It is a best practice to deploy only a single instance into a cluster.

Managed deployment

When you deploy Content Engine in an Oracle WebLogic Server managed environment, you must install and configure Content Engine on the administrative node to avoid cross-network configuration issues. Install and configure Content Engine on the administration server. Then use the administration server tools to deploy the Content Engine EAR file to the managed servers. The managed servers can be in a cluster.

Cluster deployment

Cluster deployments can be used for load balancing or to provide high availability. Refer to the *IBM FileNet P8 High Availability Technical Notice* for details on how to set up your FileNet P8 system using clusters, farms, and other high availability software and hardware.

Application server planning considerations

To prepare your application servers, you must carry out certain tasks including creating administrative accounts, specifying environmental variables, and making sure your JDBC drivers are at a supported level.

Content Engine and Application Engine are J2EE application server-based applications. (Process Engine is not.) You must install Content Engine and Application Engine in a homogeneous J2EE environment in which all of your application servers and their version numbers are identical for both components.

Assuming that a user application is required for your system and that you have not built or customized one using the IBM FileNet P8 API toolkits, you can install either of the following general user interfaces:

- Workplace XT
- Application Engine /Workplace

Even if you have your own customized application, it is a best practice to install one of these applications for testing and support purposes. The applications must use Enterprise Java Bean (EJB) transport.

If the application server where Content Engine will be deployed is running on a 32-bit JVM, it is a best practice to create no more than 75 Content Engine object stores. On a 64-bit JVM, it is a best practice to create no more than 150 Content Engine object stores.

Content Engine is a resource-intensive enterprise application. Running Content Engine and other J2EE applications on the same machine is possible but not a best practice. Other J2EE applications will compete with Content Engine for the same CPU, memory, and disk I/O resources, and increase the complexity of the installation and the risk of the deployment, because configurations will not match what has been qualified by IBM FileNet P8 Engineering.

Although you might need to host Content Engine and other applications on the same machine, it is preferable to host Content Engine on its own machine or logical partition. If an architecture requires Content Engine and a non-P8 J2EE application to be on the same machine, be sure to thoroughly test the configuration in your integration environment before deploying them into production.

See the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Authentication** for reference information about support for EJB and Web Services transports.

“WebSphere Application Server and WebLogic Server - Content Engine multi-server deployment”

In an environment of load-balanced or highly-available farmed (or clustered) application servers, you will initially install Content Engine on the Deployment Manager node (IBM WebSphere® Application Server) or the Administrator node (WebLogic Server).

WebSphere Application Server and WebLogic Server - Content Engine multi-server deployment:

In an environment of load-balanced or highly-available farmed (or clustered) application servers, you will initially install Content Engine on the Deployment Manager node (IBM WebSphere Application Server) or the Administrator node (WebLogic Server).

To install Content Engine in centrally managed servers in farms or clusters, see the *IBM FileNet P8 High Availability Technical Notice*.

Important: The base version of WebSphere Application Server that is bundled with Content Engine does not include cluster support. To configure clusters, you must have Content Engine installed on servers that are running the IBM WebSphere Application Server Network Deployment version.

Note the impact of deploying centrally managed servers not in farms or clusters. In an environment in which multiple Content Engine instances are geographically dispersed, and where each instance might have its own local Application Engine or Workplace XT server, you will install Content Engine on the Deployment Manager node (WebSphere Application Server) or the Administrator node (WebLogic Server). After installing Content Engine, you will use Configuration Manager to deploy it to other servers.

Definition of installation roles

The tasks in this guide and the rows in the Installation and Upgrade Worksheet are organized by administrative roles. Your organization might have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.

Installation administrator

- Runs IBM FileNet P8 installers during initial setup.
- Runs the Configuration Manager tool during initial setup, followed by starting IBM FileNet Enterprise Manager.
- Runs IBM FileNet P8 Upgrade programs during upgrades.
- Abbreviated as IA. Responsible for coordinating the information described in this worksheet. The information itself will require the input from the other roles.

The role of IA is usually filled by an IBM FileNet Certified Professional (FCP).

Information technology administrator

- Responsible for the networking and operating systems settings required by IBM FileNet P8.
- Responsible for performing certain security configurations.

- Abbreviated as ITA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of ITA in the Role column.

Security administrator

- Responsible for configuring the directory servers required by IBM FileNet P8 components, including Content Engine and Application Engine.
- Creates and maintains directory server user and group accounts.
- Abbreviated as SA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of SA in the Role column.

Database administrator

- Creates, configures, maintains database installations and database or table spaces.
- Responsible for creating database accounts needed by IBM FileNet P8.
- For purposes of this documentation, the database administrator is expected to have responsibilities regarding the JDBC data sources.
- Abbreviated as DBA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of DBA in the Role column.

Application server administrator

- Responsible for providing the application servers required by IBM FileNet P8.
- Responsible for application server administrative accounts.
- Abbreviated as ASA. Responsible for providing the information in the rows in the *Installation and Upgrade Worksheet* with a value of ASA in the Role column.

IBM FileNet P8 administrator

- This role designation actually refers to the administrator or administrators who perform regular maintenance of Content Engine, Process Engine Application Engine, Workplace or Workplace XT.
- The administrator who logs on to Enterprise Manager by using the `gcd_admin` account or an `object_store_admin` account is considered an IBM FileNet P8 administrator.
- Abbreviated as P8A. Responsible for providing the information in the rows of the *Installation and Upgrade Worksheet* with a value of P8A in the Role column.

E-mail Administrator

- Creates an e-mail account that will be used to configure the **Notification** tab of the Process Task Manager so that Process Engine can send e-mail notifications to end users. (Required only if you use this feature.)
- Abbreviated as EA.

Related concepts

“Using the installation and upgrade worksheet”

The Installation and Upgrade Worksheet is a Microsoft Excel spreadsheet (p8_worksheet.xls). The worksheet describes the properties and parameters required to complete IBM FileNet P8 installation, upgrade, and configuration programs, and provides a way to record the values you assign to these properties and parameters.

“IT administrator installation tasks” on page 10

The Information Technology administrator must prepare the network and operating systems, and carry out certain security configurations to prepare your environment for IBM FileNet P8 Platform.

“Security administrator installation tasks” on page 29

The Security administrator must prepare the security environment for IBM FileNet P8 Platform, including planning the security environment, configuring the directory server, and creating accounts.

“Application Server administrator installation tasks” on page 57

The Application Server Administrator must prepare the application servers for IBM FileNet P8 Platform, including planning deployment, creating administrative accounts, and configuring JDBC drivers for both Content Engine and Application Engine.

Related tasks

“Database administrator installation tasks” on page 39

The Database administrator must prepare the databases required for IBM FileNet P8, including gathering information about data sources, creating databases and database accounts, and installing client software.

Using the installation and upgrade worksheet

The Installation and Upgrade Worksheet is a Microsoft Excel spreadsheet (p8_worksheet.xls). The worksheet describes the properties and parameters required to complete IBM FileNet P8 installation, upgrade, and configuration programs, and provides a way to record the values you assign to these properties and parameters.

Administrators who are preparing the environment for installation or upgrade of IBM FileNet P8 components must use the worksheet during their preparation tasks to record the appropriate values and provide them to the Installation Administrator who runs the installation or upgrade programs.

Some of the features of the Installation and Upgrade Worksheet are:

- **Instructions:** describes the worksheet and includes a button that runs the Customize Worksheet macro.
- The two highlighted columns, **Property or Parameter** and **ENTER YOUR VALUE HERE**, provide the simplest view of the requirement. The others add identifying information and help you sort and filter the rows usefully.
- The **Role** column assigns each row to an administrator and uses the following acronyms:
 - ITA: Information Technology Administrator
 - ASA: Application Server Administrator
 - DBA: Database Administrator
 - SA: Security Administrator
 - P8A: IBM FileNet P8 Administrator
- Property definitions are contained in the column titled **Description**.

- Some rows, though not all, contain a hyperlink in the **IC help link** column. Click this hyperlink to run a query against the IBM Information Center, which opens with the Search Results pane showing the topics that contain the words in the query phrase. Browse the search results until you have enough information to be able to enter a value in the Worksheet row.

“Running the Customize Worksheet macro”

The Customize Worksheet macro lets you extract only those rows that describe your environment.

“Autofiltering and sorting the Worksheet”

There are several ways to organize the Worksheet to make finding properties and entering values easier.

Running the Customize Worksheet macro

The Customize Worksheet macro lets you extract only those rows that describe your environment.

Important: For support of the full range of built-in filter and macro features, use Microsoft Excel to view the Installation and Upgrade Worksheet file. You can use other spreadsheet programs to view the file; however, filter and macro support can vary. For example, in Calc from OpenOffice.Org, the column filters work as expected, but the Customize Worksheet button does not.

To run the Customize Worksheet macro:

1. Open the Installation and Upgrade Worksheet (p8_worksheet.xls) and click the **Instructions** worksheet (also called a tab).
2. Scroll down until you see the button representing the Customize Worksheet macro. Click the button.
3. Select the components and options that describe the environment you are preparing for IBM FileNet P8.
 - Installation or Upgrade
 - FileNet P8 Components
 - Application Server type
 - Operating system
 - Database type
 - Directory Server type
 - Number of object stores (adds new sets of rows for creating additional data sources)
 - Name of customized sheet
4. Click **OK**. The macro copies the rows that fulfill your selection criteria into a new worksheet with the name you entered. Enter the values for your environment into this new worksheet.
5. Click the name of the new worksheet at the bottom of the Excel window. Add your preparation values into this new worksheet.
6. Notice that the new worksheet has buttons at the top titled **Show Installer View** and **Show Full View**, depending on its state. The **Show Installer View** displays only those columns that you need while running installation or configuration programs.

Autofiltering and sorting the Worksheet

There are several ways to organize the Worksheet to make finding properties and entering values easier.

AutoFiltering is a quick way to display only those rows that meet a certain criteria. To use AutoFilter:

1. Make sure AutoFiltering is enabled. (Select the entire row with the column headers, then click **Data** → **Filter** → **Autofilter**.) AutoFilter arrows will appear to the right of the column labels.
2. Click the **AutoFilter** arrow in the **Installation or Configuration Program** column header and select the program you are interested in (for example, PE installer).
3. Click the **AutoFilter** arrow in the **Setup Type** column header, select **Custom**, and specify **Setup Type contains Installation**.
4. For a custom AutoFilter, click the **AutoFilter** in any column header, select **Custom**, and specify Setup Type contains "Installation".
5. To turn off AutoFiltering in a column, click the column **AutoFilter** arrow and select **(All)**.
6. To reorder rows alphabetically, do a Sort:
 - a. Click anywhere in a column, for example, Column A Role.
The only possible values in the Role column are ASA, SA, DBA, ITA, and P8A. Sorting on Role therefore groups the rows by this attribute, in alphabetic order. Several other columns also have a limited number of possible values which means they can be usefully sorted.
 - b. Click the **Sort Ascending** icon in the Excel toolbar, or use the **Data** → **Sort** menu command. The rows sort on Role.
Sorting the Worksheet reassigns row numbers. If you refer to rows by number, be aware that row numbers change if you change the sort order.

Performing the required installation preparation tasks

To efficiently carry out the required installation preparation tasks, you must assign your staff to carry out the tasks that are organized by administrative role.

Some tasks require input that results from other preparation tasks performed by other administrator roles. While performing the tasks, record results in the *Installation and Upgrade Worksheet*. See the "Using the installation and upgrade worksheet" on page 7 topic for details.

To prepare the IBM FileNet P8 environment, perform the tasks assigned to the following roles.

"IT administrator installation tasks" on page 10

The Information Technology administrator must prepare the network and operating systems, and carry out certain security configurations to prepare your environment for IBM FileNet P8 Platform.

"Security administrator installation tasks" on page 29

The Security administrator must prepare the security environment for IBM FileNet P8 Platform, including planning the security environment, configuring the directory server, and creating accounts.

"Database administrator installation tasks" on page 39

The Database administrator must prepare the databases required for IBM FileNet P8, including gathering information about data sources, creating databases and database accounts, and installing client software.

"Application Server administrator installation tasks" on page 57

The Application Server Administrator must prepare the application servers for

IBM FileNet P8 Platform, including planning deployment, creating administrative accounts, and configuring JDBC drivers for both Content Engine and Application Engine.

IT administrator installation tasks

The Information Technology administrator must prepare the network and operating systems, and carry out certain security configurations to prepare your environment for IBM FileNet P8 Platform.

- Review all rows assigned to the IT administrator (ITA) in the “Using the installation and upgrade worksheet” on page 7. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

Tip: With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following actions to quickly see only the properties assigned to a particular role:

- Click the **AutoFilter** drop-down arrow in the **Role** column header and select **ITA**.
- Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting **All**.
- If you are installing in a non-English environment, review the considerations and procedures in *IBM FileNet P8 Non-English Support Guide* before you begin your preparation tasks.

“Creating Content Engine operating system accounts”

The IT Administrator must create several operating system accounts required by FileNet P8 components during installation.

“Creating Process Engine operating system accounts” on page 13

Create new or designate existing operating system accounts for Process Engine, as shown in the following table.

“Configuring UNIX” on page 15

The IBM FileNet P8 system components require some specific configuration settings on the UNIX machines where you install them.

“Configuring the network” on page 23

You must perform certain configurations on the network before installing IBM FileNet P8 Platform.

“Preparing storage areas for object stores” on page 24

To prepare for file storage, you must prepare the locations where file stores will go, configure remote access protocol, and, if you plan to implement content-based retrieval, perform security configurations.

Creating Content Engine operating system accounts

The IT Administrator must create several operating system accounts required by FileNet P8 components during installation.

Create the following users and groups:

| User or group name | Description |
|---|---|
| Application Server Installation Administrator: (WebSphere Application Server, WebLogic Server, JBoss Application Server): <i>ce_appserver_install_user</i> | <p>An operating system user account you used to install your application server.</p> <p>Use your local machine's administrative tools to grant <i>ce_appserver_install_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> • For UNIX, <i>ce_appserver_install_user</i> must have read, write, and execute permissions to the Content Engine installation directory. <p>Later in the installation process, the <i>ce_appserver_install_user</i> performs the following tasks:</p> <ul style="list-style-type: none"> • Create and configure the application server/domain/profile for Content Engine. • Start or stop the application server when needed. • Modify the application server files or directories as needed for deploying the Content Engine application using the Configuration Manager tool. • Provide create, read and write permissions for directories on devices or drives that are used for external Content Engine file storage. <p><i>ce_appserver_install_user</i> must be a member of the <i>ce_appserver_install_group</i>.</p> |
| Application Server Installation Group: <i>ce_appserver_install_group</i> | <p>Create an operating system group account and add to it the following accounts:</p> <ul style="list-style-type: none"> • <i>ce_appserver_install_user</i> • <i>ce_install_user</i> (Windows) or <i>ce_install_user</i> (UNIX). • <i>config_mgr_user</i> <p>Use the user accounts in <i>ce_appserver_install_group</i> to do the following tasks:</p> <ul style="list-style-type: none"> • Give operating system privileges to the directories used for the Content Engine Installation and for the application server's instance/domain/profile. • Configure and deploy the Content Engine EAR files which require access to the application server's instance/domain/profile directories. • Have permissions on devices/drives to read and write that are designated for external Content Engine file storage. |

| User or group name | Description |
|---|--|
| Content Engine Installer: (UNIX): <i>ce_install_user</i> | <p>An operating system account you will use to log on to a machine to launch Content Engine installer.</p> <p>Use your UNIX administrative tools to grant this account at least the following permissions:</p> <ul style="list-style-type: none"> • Read, write, and execute permissions to the device or location where: <ul style="list-style-type: none"> – Content Engine is to be installed. – The application server instance/domain/profile has been installed. • Write permission to the directories where you will create file storage areas, index areas, and content caches. • Write permission on the <i>/tmp</i> directory. • Membership in the <i>ce_appserver_install_group</i>. <p><i>ce_install_user</i> can be the same user as the <i>ce_appserver_admin</i>.</p> |
| Configuration Manager user (WebSphere, WebLogic, JBoss): <i>config_mgr_user</i> | <p>The operating system account you will use to run Configuration Manager. At several points in <i>Plan and Prepare Your Environment for IBM FileNet P8</i> you will be instructed to grant additional permissions to <i>config_mgr_user</i>, including the following permissions:</p> <p>Add <i>config_mgr_user</i> to the <i>ce_appserver_install_group</i>.</p> <p>At several points in the <i>IBM FileNet P8 Platform Installation and Upgrade Guide</i> you will be instructed to grant additional permissions to <i>config_mgr_user</i>, including the following permissions:</p> <ul style="list-style-type: none"> • Execute permission to the Configuration Manager tool executable file, <i>configmgr.exe</i> (Windows) or <i>configmgr.sh</i> (UNIX). • Write permission to the directory where the Configuration Manager tool will create the configuration XML files. For example: <ul style="list-style-type: none"> – the directory you specify using the optional -path parameter when you run the tool – the default directory, <i>ce_install_path/tools/configurationmanager/tasks</i>, if you don't specify a path parameter • Read and write permission on the contents of the directory specified by -path parameter, or the default directory. |

| User or group name | Description |
|--|---|
| Content Engine Operating System user: <i>ce_os_user</i> | An operating system account you must log on as to create and configure the shared root directory of a file storage area or content cache area. See "Preparing storage areas for object stores" on page 24. UNIX For UNIX-based Content Engine and file storage areas, configuring security requires the use of NFS. |

Creating Process Engine operating system accounts

Create new or designate existing operating system accounts for Process Engine, as shown in the following table.

Create several operating system users and groups for Process Engine.

Create the following users and groups:

| User or group name | Description |
|--|---|
| The operating system user Process Engine uses to start and stop software (UNIX): <i>fns</i> <i>w</i> | This user is an operating system user that is used to start and stop Process Engine software. Record this value in the Installation and Upgrade Worksheet. The Process Engine installation program asks if default users will be used or if aliases for the users will be defined. If a user name other than the default FNSW is defined, indicate to the Process Engine installation program that aliases will be defined. When prompted, provide the <i>fns</i> <i>w</i> name to the installation program. Assign <i>fns</i> <i>usr</i> as the primary group for the <i>fns</i> <i>w</i> user. Assign the <i>fns</i> <i>w</i> user to the following secondary groups: <ul style="list-style-type: none"> • <i>fnadmin</i> • <i>fnop</i> |
| The password for the <i>fns</i> <i>w</i> user (UNIX): <i>fns</i> <i>w</i> <i>_password</i> | This password is for the <i>fns</i> <i>w</i> user. Record this value in the Installation and Upgrade Worksheet. |

| User or group name | Description |
|---|---|
| An operating system group for Process Engine: <i>fnop</i> | <p>This operating system group is a group account for Process Engine whose members have operator non-administrator privileges on Image Services used by Process Engine.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>Add the <i>fnsw</i> user to this group.</p> <p>If you will install Process Engine or run Process Engine Windows services as either the <i>fnsw</i> domain user or <i>pe_install_user_domain</i>, you must add that domain user to this local group.</p> <p>The Process Engine installation asks if default users and groups will be used or if aliases will be defined. If a group name other than the default <i>fnop</i> is assigned, indicate to the Process Engine installation program that aliases will be defined and provide this name to the installation program as an alias for the default group.</p> |
| An operating system group for Process Engine (UNIX): <i>fnusr</i> | <p>This operating system group is a group account for Process Engine whose members have non-administrator privileges on Process Engine files and databases.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>If you will install Process Engine or run Process Engine Windows services as either the <i>fnsw</i> domain user or <i>pe_install_user_domain</i>, you must add that domain user to this local group.</p> <p>The Process Engine installation asks if default users and groups will be used or if aliases will be defined. <i>fnusr</i> is the default group name. If a group name other than the default is assigned, indicate to the Process Engine installation program that aliases will be defined and provide this name to the installation program as an alias for the <i>fnusr</i> group. Assign the following users to the <i>fnusr</i> group:</p> <ul style="list-style-type: none"> • <i>fnsw</i> • root • Oracle user <i>oracle</i> |

| User or group name | Description |
|---|--|
| An operating system group for Process Engine (UNIX): <i>fnadmin</i> | <p>This operating system group is a group account for Process Engine whose members have all privileges on Process Engine files and databases.</p> <p>Record this value in the Installation and Upgrade Worksheet.</p> <p>The Process Engine installation asks if default users and groups will be used or if aliases will be defined. <i>fnadmin</i> is the default group name. If a group name other than the default is assigned, indicate to the Process Engine installation program that aliases will be defined and provide this name to the installation program as an alias for the <i>fnadmin</i> group. Add the following users to this group:</p> <ul style="list-style-type: none"> • <i>fnsw</i> • <i>root</i> |
| An operating system user for Process Engine (UNIX): <i>oracle</i> | <p>Create or modify an Oracle user for Process Engine.</p> <p>If the Oracle database is local to Process Engine, this user and the <i>dba</i> group might exist as a result of installing the Oracle software. If the database is remote, define the user and group on the server where Process Engine will be installed.</p> <p>Record this value in the Installation and Upgrade Worksheet to use as the Oracle UID for Process Engine installations on UNIX® with Oracle databases.</p> <p>Assign the <i>oracle</i> user to the following UNIX groups:</p> <p>Primary <i>dba</i></p> <p>Secondary <i>fnusr</i></p> |
| An operating system group for Process Engine (UNIX): <i>dba</i> | <p>Create or modify an Oracle group for Process Engine. Members of the group act as database administrators</p> <p>Record this value in the Installation and Upgrade Worksheet to use as the Oracle GID for Process Engine installations on UNIX with Oracle databases.</p> <p>Assign the <i>oracle</i> user to the <i>dba</i> group</p> |

Configuring UNIX

The IBM FileNet P8 system components require some specific configuration settings on the UNIX machines where you install them.

“Configuring UNIX for IBM FileNet P8 servers (all components)” on page 16
When configuring UNIX ensure the hosts file contents, ensure the minimum require disk and temp space, and determine your port requirements.

“Configuring Content Engine servers (all UNIX)”

For Content Engine running on a UNIX-based application server, you must use the UNIX utility program `umask` to set the default file-creation permissions mask for the JVM instance that will host Content Engine server so that the owner (the user running JVM) and the members of the owner's group have read, write, and execute access permissions, and all others have no access.

“Configuring Content Search Engine servers (all UNIX)” on page 17

The IBM FileNet P8 Content Search Engine component requires some specific configuration settings on the UNIX machines where you install it.

“Configuring Process Engine servers (all UNIX)” on page 17

The following operating system prerequisites apply to all UNIX-based IBM FileNet P8 Process Engine servers.

“Configuring Process Engine servers (Solaris)” on page 19

The operating system prerequisites in this subsection pertain to Solaris-based Process Engine servers.

“Configuring Application Engine (Solaris)” on page 23

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons.

Configuring UNIX for IBM FileNet P8 servers (all components):

When configuring UNIX ensure the hosts file contents, ensure the minimum require disk and temp space, and determine your port requirements.

To configure UNIX for FileNet P8 Servers:

1. Ensure hosts file contents. On each UNIX-based FileNet P8 server that does not use DNS (Domain Name Service) or NIS (Network Information Service), the `/etc/hosts` file must contain the name and Internet Protocol (IP) address of all servers it will communicate with, including the remote database server, if applicable. Process Engine has additional requirements for hosts file entries. See “Configuring the `/etc/hosts` file” on page 18 for additional information.
2. Ensure minimum required disk space and temp space for installation. See the IBM FileNet P8 4.5 Hardware and Software Requirements guide.
3. Determine port requirements. Consult with the application server, database, and FileNet P8 administrators to determine port requirements for all the servers in your installation environment.
4. If you intend to install FileNet P8 components interactively, ensure that each UNIX-based FileNet P8 server uses the X Window system. And if you are going to install components from a remote machine, verify that the remote machine has an X Window terminal emulator.

Configuring Content Engine servers (all UNIX):

For Content Engine running on a UNIX-based application server, you must use the UNIX utility program `umask` to set the default file-creation permissions mask for the JVM instance that will host Content Engine server so that the owner (the user running JVM) and the members of the owner's group have read, write, and execute access permissions, and all others have no access.

```
umask u=rwx,g=rwx,o=
```

This mask setting ensures that the access permissions on files and directories created by Content Engine server are identical to those you must specify when creating file storage areas on UNIX file servers.

This umask setting is also required for the user running Content Engine setup (*ce_install_user*). The umask must be in the *.profile* file for the user running JVM and also for *ce_install_user*.

Configuring Content Search Engine servers (all UNIX):

The IBM FileNet P8 Content Search Engine component requires some specific configuration settings on the UNIX machines where you install it.

Important: Ensure that the file compression utilities, *gzip* and *gunzip*, are at version 1.3 (or higher) on any UNIX machine where you will install Autonomy K2 locale files.

“Setting up Java SDK for Content Search Engine”

Content Search Engine requires Java SDK on UNIX machines where you intend to install additional locales (other than English).

Setting up Java SDK for Content Search Engine:

Content Search Engine requires Java SDK on UNIX machines where you intend to install additional locales (other than English).

To set up Java SDK:

On each UNIX machine where you intend to install an additional locale (other than English), complete the following substeps:

1. Install the Java 1.4 SDK.
2. Set the PATH environment variable to *install_path/Java14/bin*, where *install_path* is the path to where you installed Java 1.4 SDK.

Configuring Process Engine servers (all UNIX):

The following operating system prerequisites apply to all UNIX-based IBM FileNet P8 Process Engine servers.

“Configuring UNIX servers for Process Engine”

Perform the following prerequisite tasks in any order.

“Configuring the */etc/hosts* file” on page 18

Information related to Process Engine IP address, server name and optional load balancer must be entered into either the server's DNS table or the hosts file.

Configuring UNIX servers for Process Engine:

Perform the following prerequisite tasks in any order.

(Solaris only): Solaris volume management software might use port 32776. This is the default for the Process Engine Communication Port (IOR port).

To configure UNIX servers:

- Ensure minimum */tmp* size. The */tmp* directory must have 510 MB free.
- Save the following files for the root user.
 - *.cshrc*
 - *.Xdefaults*
 - *.Xresources*
 - *.dbxinit*

- .dtprofile
 - .env
 - .login
 - .mwmrc
 - .xinitrc
 - .profile
- Before installing Process Engine verify that your operating system is set up with a correctly configured volume manager. You can use the volume manager provided with the operating systems or an equivalent Veritas volume manager.
 - Process Engine requires the presence of several partitions.

| Volume Name | Mount Point | Minimum Size | User | Group | Mode |
|-------------|-------------------------|--------------|------|-------|------|
| fnsw | /fnsw file system | 2 GB | fnsw | fnusr | 775 |
| local | /fnsw/local file system | 10 GB | fnsw | fnusr | 775 |

Configuring the /etc/hosts file:

Information related to Process Engine IP address, server name and optional load balancer must be entered into either the server's DNS table or the hosts file.

For non-farmed configurations, the information can be in either the DNS table or the hosts file on the server. For farmed configurations, the information must be entered into the hosts file. In a farmed environment, entries must exist for every Process Engine server in the farm.

Process Engine can be configured to run with IPv4, IPv6, or a dual stack network.

Entries must use the following format for each Process Engine server. The load balancer name must also be associated with the appropriate server in a farmed configuration.

IP_addr hostname load_balancer_name

IP_addr

The IP address of the Process Engine server.

hostname

The corresponding host name, also referred to as the network name. This name must be < 32 characters and cannot contain any dots. Record this value in the Installation and Upgrade Worksheet as the network name and provide it to the Process Engine installation program.

load_balancer_name

The name of the load balancer in a farmed configuration.

In the following example there are IPv4 and IPv6 addresses and cmepwin11 is the hostname (network name). If the server is configured for a dual stack network, both IPv4 and IPv6 entries must exist for the same hostname.

```
127.0.0.1      localhost
::1           localhost
2007:0:0:0:9:181:124:192 cmepwin11
#fe80::fda2:48f4:7106:2f19%10 cmepwin11
123.45.6.78   cmepwin11
```

Configuring Process Engine servers (Solaris):

The operating system prerequisites in this subsection pertain to Solaris-based Process Engine servers.

“Enabling ports”

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons.

“Verifying national language character set and time settings”

Verify the national language character set and time settings on a Solaris-based Process Engine server.

“Increasing the operating system kernel limits (Solaris 9)” on page 20

Increase or verify Solaris 9 operating system kernel limits before installing Process Engine software.

“Increasing the operating system kernel limits (Solaris 10)” on page 21

Increase or verify Solaris 10 operating system kernel limits before installing Process Engine software.

Enabling ports:

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons.

By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet P8 uses several ports higher than 32768.

To use these ports on Solaris-based systems, you must first enable the ports by setting the smallest anon port to 32778. By doing so, the ports used by Solaris communication daemons will be 32778 or greater, leaving 32777 available for IBM FileNet P8 use.

The Solaris platform provides several different tools, such as the `netstat` command, to determine if a port is in use.

1. Determine the current `tcp_smallest_anon_port` setting. From a command prompt, enter the following command: `ndd -get /dev/tcp tcp_smallest_anon_port`.
2. Enable port 32777. If the port returned in the step above is less than 32778, you must enable port 32777.

Solaris 9

Edit the `/etc/rc2.d/S69inet` file. Add the following line before the `exit 0` entry at the bottom of the file:

```
ndd -set /dev/tcp tcp_smallest_anon_port 32778
```

Solaris 10

Edit the `/lib/svc/method/net-init` file. Add the following line before the `exit 0` entry at the bottom of the file:

```
ndd -set /dev/tcp tcp_smallest_anon_port 32778
```

3. Reboot the Process Engineserver to force the release of ports required by Process Engine that might be in use by the operating system. Failure to reboot after these changes are made can result in the port being unavailable, generating OpenSocket errors.

Verifying national language character set and time settings:

Verify the national language character set and time settings on a Solaris-based Process Engine server.

To verify the national language character set and time settings:

1. The default time mask varies on UNIX depending on the *LANG* and *LC_TIME* environment settings. Verify the current *LC_TIME* settings by entering: `locale -k t_fmt` The result might appear similar to this:

```
t_fmt=%r
```

2. The default mask must not be `%r`. To change to a default mask that can be used with *NLT*, reset the *LC_TIME* environment to *C*, then run the `t_fmt=%r` command again to verify the change.
3. Modify the default shell environment to use the *C* time format.
4. Change the `/etc/profile` for the entire system or change `.profile` files for each user that runs *sh* or *ksh* to include the following lines:

```
LC_TIME=C
export LC_TIME
```

5. Verify the current *LANG* settings by entering `locale` at the shell prompt. This example shows the U.S. character set, *ISO 8859-1*. Be sure it is consistent with the database character set unless your database character set is *AL32UTF8* (Unicode).

```
LANG=en_US.ISO8859-1
LC_CTYPE=en_US.ISO8859-1
LC_NUMERIC=en_US.ISO8859-1
LC_TIME=en_US.ISO8859-1
LC_COLLATE=en_US.ISO8859-1
LC_MONETARY=en_US.ISO8859-1
LC_MESSAGES=en_US.ISO8859-1
LC_ALL=
```

Increasing the operating system kernel limits (Solaris 9):

Increase or verify Solaris 9 operating system kernel limits before installing Process Engine software.

To increase the operating system kernel limit:

1. Make a copy of the system file (with a new name). Log on as root, and enter a command similar to the following: `cp /etc/system /etc/system.save`
2. Edit the `/etc/system` file, using your preferred editor (for example, *vi*):
`vi /etc/system`
3. Ensure that the following parameters are listed and are set to at least the values shown.

```
set semsys:seminfo_semmap=50
set semsys:seminfo_semmni=2000
set semsys:seminfo_semmns=2000
set semsys:seminfo_semmnu=500
set semsys:seminfo_semmsl=512
set semsys:seminfo_semopm=256
set semsys:seminfo_semume=500
set semsys:seminfo_sevmx=32767
set semsys:seminfo_semaem=16384
set shmsys:shminfo_shmmax=4294967295*
set shmsys:shminfo_shmmi=0
set shmsys:shminfo_shmmni=2000
set shmsys:shminfo_shmseg=100
set msgsys:msginfo_msgmni=2048
set max_nprocs=1000
set fnsod:sod_Debug=0
set rlim_fd_max=1024
set rlim_fd_cur=256
noexec_user_stack=1
```

Increasing the operating system kernel limits (Solaris 10):

Increase or verify Solaris 10 operating system kernel limits before installing Process Engine software.

The values of the parameters in the project file override the values of the parameters in the `/etc/system` file. Make sure that the values of the parameters in the project file are the same as or greater than the values of the parameters in the `/etc/system` file.

Solaris 10 provides the Resource Control Facility feature for setting kernel parameters. Refer to the Solaris 10 documentation on the Sun Web site for information about this feature. Kernel parameter changes should be made to the Solaris Resource Control Facility whether or not you are using zones in Solaris 10. In a configuration with zones, make these changes in the global zone.

To increase the operating system kernel limits:

1. Log on as `fns`.
2. At a shell prompt enter `prctl $$`.
3. Verify the values are set to at least the minimum values indicated for Process Engine.

| Kernel Parameter | Solaris 10 default | Solaris 10 for Process Engine |
|--|-------------------------------------|-------------------------------|
| <code>max-sem-nsems</code> | 512 | 1024 |
| <code>max-sem-ops</code> | 512 | 1024 |
| <code>max-sem-ids</code> | 128 | 2048 |
| <code>max-shm-ids</code> | 128 | 2048 |
| <code>max-msg-ids</code> | 128 | 2048 |
| <code>max-shm-memory</code> | 1/4 Physical memory | 1/4 Physical memory |
| <code>max_nprocs</code> | 10 + (16 x <code>maxusers</code>)* | default |
| <code>Rlim_fd_max</code> | 65536 | 65536 |
| <code>Rlim_fd_cur</code> | 256 | 256 |
| <code>Noexec_user_stack</code> | 0 | 1 |
| <code>max-msg-qbytes</code> | 65536 | 65536 |
| <code>Msgtql</code> (process.max-msg-messages) | 8192 | 65536 |

* `maxusers` - 2048 or server physical memory (in MB), whichever is less.

4. Logged on as the root user, copy the project and the system file.

```
cp /etc/project /etc/project.save  
mv /etc/system /etc/system.save
```
5. View the current settings for the `fns` user, and the Oracle user. For the FileNet software user, enter:

```
su - fns  
projects -l
```

For Oracle user, enter:

```
su - oracle  
projects -l
```

6. Return to the user with root privileges.

7. If a project does not exist for the fnsw user, enter the following command to add one:

```
projadd -U fnsw -K  
"project.max-shm-memory=(priv,7G,deny)" user.fnsw
```

8. If a project does exist for the fnsw user, enter the following command:

```
projmod -sK  
"project.max-shm-memory=(priv,7G,deny)" user.fnsw
```

Update other kernel parameters required for the fnsw user in /etc/project, for example:

```
projmod -sK "project.max-shm-ids=(priv,2048,deny)"  
user.fnsw  
projmod -sK "project.max-msg-ids=(priv,2048,deny)" user.fnsw  
projmod -sK "project.max-sem-ids=(priv,2048,deny)" user.fnsw  
projmod -sK "process.max-msg-messages=(priv,65536,deny)"  
user.fnsw  
projmod -sK "process.max-msg-qbytes=(priv,65536,deny)"  
user.fnsw  
projadd -U fnsw -K "project.max-shm-memory=(priv,7G,deny)"  
user.fnsw  
projmod -sK "project.max-sem-nsems=(priv,1024,deny)"  
user.fnsw  
projmod -sK "project.max-sem-ops=(priv,1024,deny)" user.fnsw  
projmod -sK  
"project.max-device-locked-memory=(priv,3.9G,deny)" user.fnsw  
projmod -sK "project.max-shm-ids=(priv,2k,deny)" user.fnsw  
projmod -sK "process.max-msg-messages=(priv,1.02k,deny)"  
user.fnsw  
projmod -sK "process.max-sem-ops=(priv,1k,deny)" user.fnsw  
projmod -sK "process.max-sem-nsems=(priv,1.02k,deny)"  
user.fnsw
```

9. If a project does not exist for the RDBMS software user, such as oracle, enter the following command to create one:

```
projadd -U oracle -K  
"project.max-shm-memory=(priv,7G,deny)" user.oracle
```

10. If a project does exist for the RDBMS user, such as oracle, enter the following command:

```
projmod -sK  
"project.max-shm-memory=(priv,7G,deny)" user.oracle
```

11. Update the following kernel parameters. These changes are not necessary if the value is already set to a higher value.

```
projmod -sK  
"project.max-shm-memory=(priv,7G,deny)" user.root  
projmod -sK "project.max-shm-memory=(priv,7G,deny)" system  
projmod -sK "project.max-shm-memory=(priv,7G,deny)"  
default
```

12. Make a copy of the system file (with a new name):

```
cp /etc/system /etc/system.save
```

13. Edit the /etc/system file, using your preferred editor (for example, vi):

```
vi /etc/system
```

Add the following value to the file:

```
set noexec_user_stack=1
```

14. Exit the file and save your changes.

15. Verify the new active kernel values for both FileNet and Oracle users match the projects -l output you saw in 5 on page 21.

For fnsw user, enter:

```
su - fnsw
prctl $$
For Oracle user, enter:
su - oracle
prctl $$
```

16. Verify the `/etc/system` file by entering:
`more /etc/system`

Configuring Application Engine (Solaris):

When Solaris starts up, it takes the first several ports, called anon ports, to use for its communication daemons.

By default, the maximum `tcp_smallest_anon_port` is 32768. IBM FileNet P8 uses several ports higher than 32768.

To use these ports on Solaris-based systems, you must first enable the ports by setting the smallest anon port to 32778. By doing so, the ports used by Solaris communication daemons will be 32778 or greater, leaving 32777 available for IBM FileNet P8 use.

The Solaris platform provides several different tools, such as the `netstat` command, to determine if a port is in use.

1. Determine the current `tcp_smallest_anon_port` setting. From a command prompt, enter the following command: `ndd -get /dev/tcp tcp_smallest_anon_port`
2. Enable port 32777. If the port returned in the step above is less than 32778, you must enable port 32777.

Solaris 9

Edit the `/etc/rc2.d/S69inet` file. Add the following line before the exit 0 entry at the bottom of the file:

```
ndd -set /dev/tcp tcp_smallest_anon_port 32778
```

Solaris 10

Edit the `/lib/svc/method/net-init` file. Add the following line before the exit 0 entry at the bottom of the file:

```
ndd -set /dev/tcp tcp_smallest_anon_port 32778
```

3. Reboot the Application Engine server to force the release of ports required by the Application Engine that might be in use by the operating system. Failure to reboot after these changes are made can result in the port being unavailable, generating OpenSocket errors.

Configuring the network

You must perform certain configurations on the network before installing IBM FileNet P8 Platform.

“Prerequisites to configuring your network”

Perform the following prerequisite tasks in any order.

“Synchronizing time and date” on page 24

System users will experience a variety of problems if one or more servers are not synchronized with the rest of the system.

Prerequisites to configuring your network:

Perform the following prerequisite tasks in any order.

- Assign all IBM FileNet P8 servers a static IP address.

- Ensure TCP/IP settings. Verify TCP/IP configuration settings on all UNIX and Windows servers and Enterprise Manager clients intended for IBM FileNet P8 so that they can all communicate with one another.
- Ensure availability of required port numbers. Several port numbers are required by the various IBM FileNet P8 components.

Synchronizing time and date:

System users will experience a variety of problems if one or more servers are not synchronized with the rest of the system.

The Process Engine database server (the machine that hosts the database used by Process Engine) is considered the master time keeper; the UTC time of that machine is considered the correct time. The server hosting the Process Engine API and the server hosting Content Engine must have the UTC time set.

- To change the time on the machine hosting Process Engine, you must stop the server. In a farmed Process Engine system, if you want to change the time of one of the servers in the farm, you must stop only that server.
- To change the time in the machine hosting the Process Engine API, be sure it is not connected to any Process Engine system. If the API is connected to a Process Engine server, and you change the time, you will experience authentication errors, and you might need to log on again.
- If your Content Engine server is being used with a Process Engine server, and you change the time on the Content Engine server, you will experience authentication errors in Process Engine and you might need to log on again.

Preparing storage areas for object stores

To prepare for file storage, you must prepare the locations where file stores will go, configure remote access protocol, and, if you plan to implement content-based retrieval, perform security configurations.

An object store can have up to three types of storage areas for the content of documents and business objects:

file storage area

A file storage area stores content in a network-accessible directory. The path name to this directory specifies the location of the file storage area.

For information about file storage areas, see the IBM FileNet P8 help topic **System Administration** → **Content Engine Administration** → **Content storage** → **File storage areas**.

fixed storage area

A fixed storage area is a file storage area on a large-capacity, (possibly) write-once, fixed content device.

For information about fixed storage areas, see the IBM FileNet P8 help topic **System Administration** → **Content Engine Administration** → **Content storage** → **File storage areas**.

database storage area

A database storage area stores content as binary large objects (BLOBs) in a database.

A *file storage area* refers only to a network-accessible directory that is not on a fixed content device. The names of file storage areas and database storage areas must be unique within an object store. File storage areas on encrypted NTFS devices are not

supported. However, the Decru, Vormetrics, and IBM Encryption Expert hardware-based encryption solutions are supported.

By default, the Create Object Store wizard creates a database storage area. If your object stores will use database storage areas only, you can skip the rest of this task, provided that one of the following conditions is met:

- Your database type is non-DB2.
- Your database type is DB2 and your database storage areas will not contain large content elements (larger than 300 MB, for example).

Besides creating a database storage area, the Create Object Store wizard allows you to create an initial file storage area or initial fixed storage area. But the wizard requires that you first do at least one of the following, depending on the type of storage areas you want for your object stores:

- For fixed storage areas, create at least one fixed content device (typically via Enterprise Manager). Multiple fixed storage areas can share the same fixed content device, or a fixed storage area can have its own fixed content device.

If the content of all your object stores will be in fixed storage areas only, create your fixed content devices now, and skip the rest of this topic.

To create a fixed content device, refer to the procedures in IBM FileNet P8 help topic **System Administration** → **Content Engine Administration** → **Content storage** → **File storage areas**.

- For file storage areas, prepare locations on one or more file servers (which usually are not a machine where you installed Content Engine), as shown in the remainder of this task.

“Configuring file servers for file storage areas”

You must configure file servers for the initial file storage areas of the object stores to be created, and for additional file storage areas of existing object stores.

“Users and groups” on page 26

The following table shows the operating system users and groups on the machine where Content Engine is to be deployed that are involved in securing file storage areas. These users and groups must be defined in the directory service that the operating system uses to authenticate users, which is not necessarily the same directory service that Content Engine Server uses.

“Configuring the remote access protocol on the client machine” on page 28

When configuring the remote file access protocol (NFS or CIFS), the client machine is the one where Content Engine Server or Content Search Engine are running. Configuring the remote access protocol (NFS or CIFS) means designating a directory (where content is be stored) so that it appears to be on the a local file system of the client machine.

“Preparing storage areas to support content search functionality” on page 28

To prepare your system for index area creation, each file storage area that will be full-text indexed must be accessible by both the Content Engine and the Autonomy K2 server that will perform the full-text indexing.

Configuring file servers for file storage areas:

You must configure file servers for the initial file storage areas of the object stores to be created, and for additional file storage areas of existing object stores.

Refer to the *IBM FileNet P8 Hardware and Software Requirements* for currently supported operating systems for file servers.

Configuring a file server for file storage areas involves the following general steps, which are described in more detail in the procedures later in this task.

To configure file servers for file storage areas:

1. Create or designate an existing top-level directory on the file server where file storage areas will reside.
2. Secure the directory so only Content Engine Server and Content Search Engine can access it.
3. Expose the directory via the remote file access protocol that applies to the operating system of the file server.
4. (Best practice) Under the top-level directory, create a subdirectory for each file storage area you intend to create. If you decide to put a file storage area directly within a top-level directory, rather than in a subdirectory, and you later decide to create an additional file storage area on this file server, you will have to create another top-level directory for it, because you will not be able to use the previously created top-level directory.

“Remote file access protocols”

The supported remote file access protocols between Content Engine and a file server are: Common Internet File System (CIFS), Network File System (NFS), and Distributed File System (DFS). DFS is supported if you are using it to manage a file storage area; however, the replication feature of DFS is not supported.

Remote file access protocols:

The supported remote file access protocols between Content Engine and a file server are: Common Internet File System (CIFS), Network File System (NFS), and Distributed File System (DFS). DFS is supported if you are using it to manage a file storage area; however, the replication feature of DFS is not supported.

The communication method between the Content Engine machine and the file server depends on the operating systems running on the two machines, as shown in the following table:

| Content Engine Operating System | File Server Operating System | File Access Protocol |
|--|-------------------------------------|-----------------------------|
| Windows 2003 | Windows 2003 | CIFS |
| UNIX | UNIX | NFS |
| UNIX | Windows 2003 | NFS |

Install a UPS power supply backup system on each file server to enable graceful shutdown. Loss or corruption of data will result if a file server does not shut down gracefully.

Users and groups:

The following table shows the operating system users and groups on the machine where Content Engine is to be deployed that are involved in securing file storage areas. These users and groups must be defined in the directory service that the operating system uses to authenticate users, which is not necessarily the same directory service that Content Engine Server uses.

The user and group account variables in this table are placeholders for the actual account names that you designate.

| Users and Groups | Role |
|---|---|
| Content Engine operating system user (ce_os_user) | The user under which Content Engine Server runs (typically, the user that starts Content Engine Server). |
| K2 operating system user (k2_os_user) | The user under which Content Search Engine runs (typically, the user that starts Content Search Engine). |
| Content Engine operating system group (ce_os_group) | The group that contains: <ul style="list-style-type: none"> • Content Engine operating system user • K2 operating system user |

For details on file storage area security, see the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Authorization** → **Storage area security**.

“Configuring a UNIX-based file server”

You need to create a directory and specify permissions for the Content Engine operating system user before you can create a storage area.

Configuring a UNIX-based file server:

You need to create a directory and specify permissions for the Content Engine operating system user before you can create a storage area.

To configure a UNIX-based file server:

1. Log on to the UNIX file server as a user with read/write access to the device where you want to create a storage area.
2. Create or designate a directory for the first storage area where content will be stored (as in, fsa1). For example:

```
$ mkdir /opt/filenet/file_stores/fsa1
```

3. Set the Content Engine operating system user as the owner of fsa1 and give group access permission to the Content Engine operating system group. For example:

```
chown ce_os_user:ce_os_group fsa1
```

Tip: The UID (user ID) for *ce_os_user* and the GID (group ID) for *ce_os_group* on the file server must match the UID and GID for the same user and group on the machine where Content Engine and Content Search Engine are running. This will normally be true if all machines use the same directory service, but they might be different.

4. Change the permissions on fsa1 so that *ce_os_user* and *ce_os_group* both have read/write/execute privileges and all other users have no privileges:

```
chmod 0770 fsa1
```

5. Via NFS, export fsa1. Alternatively, if the file server will host more than one file storage area, export the parent directory. In the latter case, for example, export /opt/filenet/file_stores, rather than /opt/filenet/file_stores/fsa1, and then create a separate subdirectory to serve as the root of each file storage area.

Tip: It is a best practice to restrict trusted hosts to just those on which an instance of Content Engine Server or Content Search Engine is executing. Root access should also be restricted. Refer to the UNIX administrator manual for details on exporting files in NFS.

Configuring the remote access protocol on the client machine:

When configuring the remote file access protocol (NFS or CIFS), the client machine is the one where Content Engine Server or Content Search Engine are running. Configuring the remote access protocol (NFS or CIFS) means designating a directory (where content is to be stored) so that it appears to be on the local file system of the client machine.

To configure remote access protocol:

To configure UNIX-based Content Engine Server to talk to a UNIX or Windows file server via NFS:

1. On the application server where you are going to deploy Content Engine Server, log on as the user who launched the application server.
2. Mount the exported NFS file system (from step 5 on page 27 of “Configuring a UNIX-based file server” on page 27) onto a local directory on the Content Engine machine. The mount point must be in the same location on all machines where Content Engine Server and Content Search Engine are going to be installed in the local file system.

For example:

```
mount -t nfs filesrv :/opt/filenet/file_stores/home/filenet/file_stores
```

where *filesrv* is the host name of Content Engine machine.

In this example, all Content Engine Server machines (including machines that are part of the same server farm or cluster) must mount the remote file system at */home/filenet/file_stores*.

Preparing storage areas to support content search functionality:

To prepare your system for index area creation, each file storage area that will be full-text indexed must be accessible by both the Content Engine and the Autonomy K2 server that will perform the full-text indexing.

Content Engine Operating System User (*ce_os_user*) must have permissions to the following directories:

- File storage area directories (read and write)
- Collections directory (read and write)
- Collections temp directory (read and write)

K2 Operating System User (*k2_os_user*) must have permissions to the following directories:

- File storage area directories (read)
- Collections directory (read and write)
- Collections Temp directory (read and write)

If Content Engine and Autonomy K2 are installed on separate machines, then the file store must be mounted and accessible to both the Content Engine and the Autonomy K2 servers. For Windows, the file store path must be a UNC path name

accessible to the Autonomy K2 machine. For UNIX, the file store must be NFS-mounted on the Autonomy K2 machine with the same name as it appears on all other machines.

Related tasks

“Creating Content Engine operating system accounts” on page 10

The IT Administrator must create several operating system accounts required by FileNet P8 components during installation.

“Creating Content Search Engine accounts” on page 38

If you are installing Content Search Engine, you must create new or designate existing Autonomy K2 security accounts.

Security administrator installation tasks

The Security administrator must prepare the security environment for IBM FileNet P8 Platform, including planning the security environment, configuring the directory server, and creating accounts.

Review all rows assigned to the Security administrator (SA) in the Installation and Upgrade Worksheet. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following actions to quickly see only the properties assigned to a particular Role:

- Click the **AutoFilter** drop-down arrow in the **Role** column header and select **SA**.
- Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting **All**.

“Security planning considerations”

Information in this section is provided to assist in the security planning process but is not a complete description of any security feature or level of support.

“Configuring directory server” on page 31

The Security administrator must perform certain configurations on the directory server that will provide the authentication repository for your IBM FileNet P8 system.

“Specifying IBM FileNet P8 accounts” on page 33

Included in this topic you will find a set of tables that describe all the accounts you must specify to set up IBM FileNet P8 components.

Security planning considerations

Information in this section is provided to assist in the security planning process but is not a complete description of any security feature or level of support.

For complete information about IBM FileNet P8 security, consult the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security**. The Security Help sections that are especially applicable to installation are Authentication, Directory Service Providers, and Users and Groups.

Authentication and authorization are separate processes.

Authentication (logon security) is separate from authorization (object and process security). You must configure your JAAS login on the Content Engine application server so that any user or group that can successfully log on to IBM FileNet P8 resources can also be authorized to work within IBM FileNet P8 interfaces, using the Content Engine directory service provider connection.

Configuration Manager captures configuration information to create your application server authentication provider; or you can use an authentication provider that already exists on the application server. Immediately following the initial Content Engine deployment, you will use Enterprise Manager to configure the Content Engine authorization by creating a Directory Configuration.

Logins are done through JAAS.

IBM FileNet P8 uses Java Authentication and Authorization Service (JAAS) for authentication, which is a process that occurs between a J2EE client application, a J2EE application server, and one or more JAAS login modules. This process does not involve any IBM FileNet P8 code.

IBM FileNet P8 Platform uses JAAS for authentication only, not for authorization on stored objects. Also, it does not support Java Security Manager.

Determine single sign-on (SSO) requirements.

Content Engine ability to use JAAS-based authentication means that if a single sign-on (SSO) provider writes a JAAS LoginModule for a supported application server, then clients of IBM FileNet P8 applications hosted in that application server can use that SSO solution. Where necessary, this Information Center describes SSO configurations that you must make, for example to the `web.xml` file, but it does not provide specific instructions for installing or configuring your SSO provider. See the IBM FileNet P8 help topic **FileNet P8 Administration → Enterprise-wide Administration → FileNet P8 Security → Authentication → Single sign-on integrations via JAAS**. Also see *Single Sign-On Solutions for IBM FileNet P8* at ibm.com/redbooks for configuration information.

Determine Kerberos applicability.

You can use Kerberos for SSO authentication between IBM FileNet Enterprise Manager and Content Engine, provided you use Windows Active Directory as the directory server. See the IBM FileNet P8 help topic **FileNet P8 Administration → Enterprise-wide Administration → FileNet P8 Security → Authentication → Kerberos for Content Engine**.

For information, see the IBM FileNet P8 help topic **FileNet P8 Administration → Enterprise-wide Administration → FileNet P8 Security → Authentication**.

Decide how many authentication realms you require.

At least one authentication realm is required, which you create during an initial installation by running Configuration Manager's Configure LDAP task. For an explanation of how to configure multiple realms, see the IBM FileNet P8 help topic **FileNet P8 Administration → Enterprise-wide Administration → FileNet P8 Security → How to → Configure multiple realms**.

Make sure that you have a directory service provider in place.

Authentication in IBM FileNet P8 Platform is provided by the following supported directory servers:

- Sun Java System Directory Server

You can find detailed reference information in the IBM FileNet P8 help topic **FileNet P8 Administration → Enterprise-wide Administration → FileNet P8 Security → Directory Service Providers**.

IBM FileNet P8 supports only homogenous directory server environments. In other words, a single IBM FileNet P8 domain can be configured to use only one of the supported directory servers.

Understand the users and groups required for IBM FileNet P8.

All general administrative users and groups needing access to IBM FileNet P8-based applications must reside in one of the supported directory servers. This *Plan and Prepare Your Environment for IBM FileNet P8* provides instructions for creating the administrative accounts required for installation and initial configuration. For detailed reference information, see the IBM FileNet P8 help topic **FileNet P8 Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Users and groups**.

Process Engine delegates authentication to Content Engine.

Start with the IBM FileNet P8 Platform 4.0.0 release, Process Engine no longer has a direct connection to a directory server for authentication purposes, as it did in earlier releases. Instead, it delegates authentication tasks to Content Engine.

(WebLogic only) Any WebLogic authentication provider must be dedicated to IBM FileNet P8.

For performance reasons, do not share any authentication provider used by WebLogic for deployed IBM FileNet P8 components with applications used for other purposes.

You can configure Content Engine to use email or UPN for login

You can assign the directory server's email attribute or, for Active Directory, the userPrincipalName (UPN) to be the user short name used for login. Instructions in the *IBM FileNet P8 Platform Installation and Upgrade Guide* provide a link to a procedure that explains how to do this.

Configuring directory server

The Security administrator must perform certain configurations on the directory server that will provide the authentication repository for your IBM FileNet P8 system.

“Configuring Sun Java System Directory Server”

You can configure Sun Java System Directory Server on a Windows or UNIX server to be the directory service for IBM FileNet P8 Platform.

Configuring Sun Java System Directory Server:

You can configure Sun Java System Directory Server on a Windows or UNIX server to be the directory service for IBM FileNet P8 Platform.

On Windows servers, Sun Java System Directory Server should be installed on an NTFS hard drive partition.

If there are more than 2,000 users in the Directory Server, you must increase the resource limits to correctly display users in IBM FileNet P8. IBM recommends setting this limit to -1 (unlimited). You can either set this limit for the entire LDAP server or for the individual IBM FileNet P8 users.

Server Side Sorting (SSS) must be enabled. This is because Process Engine and other IBM FileNet P8 components call on Content Engine to perform searches using a sorted paging mechanism. Without SSS, you will experience errors such as in Workplace when retrieving a document, when trying to open Process Designer

via Workplace, or when starting a Connection Point on Application Engine. Note that SSS is normally enabled by default but is sometimes disabled due to concerns with performance.

For a complete list of IBM FileNet P8-supported Sun Java System Directory Server features, refer to the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Directory Service Providers** → **Sun Java System Directory Server**.

“Setting the resource limits for the entire Directory Server (v 5.2)”

User resource limits take precedence over server resource limits. Existing users who have a value specified for resource limits will not be affected by the changes made in the following steps.

“Setting the resource limits for individual IBM FileNet P8 users”

You must perform these steps any time you add IBM FileNet P8 users.

Setting the resource limits for the entire Directory Server (v 5.2):

User resource limits take precedence over server resource limits. Existing users who have a value specified for resource limits will not be affected by the changes made in the following steps.

To set the resource limits:

1. On the server where Sun Java System Directory Server is installed, log in with an account that has rights to modify the Sun Java System Directory Server environment.
2. Run the Sun Java System Directory Server console and log in.
3. Expand the **Domain** → **Server Group** containers and select your *Directory Server*.
4. Right-click and select **Open**.
5. Select the **Configuration** tab.
6. Select the **Performance** container.
7. Select the **Client Control** tab.
8. For the **LDAP** group box, ensure that **Size limit** and **Look-through limit** are both set to **Unlimited**.
9. If changes were made, click **Save**.
10. Select the **Tasks** tab and **Restart the Directory Server** if changes were made.

Setting the resource limits for individual IBM FileNet P8 users:

You must perform these steps any time you add IBM FileNet P8 users.

To set the resource limits for individual users:

1. From the Sun Java System Directory Server console, expand the **Domain** → **Server Group** containers and select your Directory Server. Then click **Open**.
2. Select the **Directory** tab.
3. From the left pane, select the **Object** (OU, etc.) that contains the user(s) you want to change.
4. For each IBM FileNet P8 user whose limit you want to change, complete the following steps:
 - a. From the right pane, double-click on the user name.
 - b. Select **Properties**.

- c. On the left pane of the **Properties** dialog box, select **Account**.
 - d. Enter -1 in the **Look through limit** and **size limit** fields.
 - e. Click **OK**.
5. Restart the Directory Server.

Specifying IBM FileNet P8 accounts

Included in this topic you will find a set of tables that describe all the accounts you must specify to set up IBM FileNet P8 components.

This task assumes that you have completed the tasks required for your directory server in “Configuring directory server” on page 31.

Although these accounts are collected here under the Security administrator tasks because of their relationship to overall system security concerns, you will notice mention of other administrators who are likely to have the actual responsibility to create the accounts and put them to use.

The following procedures direct you to create or designate the accounts needed to install and configure IBM FileNet P8. For a complete list of the user and group roles, accounts, and responsibilities required to install, configure, and maintain an IBM FileNet P8 system, see the IBM FileNet P8 help topic **System Administration** → **Enterprise-wide Administration** → **FileNet P8 Security** → **Users and groups**.

Accounts are referred to in documentation in the following ways:

- By a display name; for example, *Database User Name*. An account's display name is how the IBM FileNet P8 user interface, such as an installation program or dialog box, refers to the account. Many accounts have both a display name and a variable.
- By a variable designator; for example *ce_db_user*, using lower-cased italics and underscores. The variable is intended to show that you must designate your own account to act in the role described by the variable. Accounts that do not appear in an interface or configuration file will have only a variable designator.

If you see a reference to an account that you do not understand, search the *Plan and Prepare Your Environment for IBM FileNet P8* guide and find the account table that defines it.

“Creating Content Engine directory server accounts”

Create new or designate existing directory server installation accounts for Content Engine.

“Creating Process Engine directory server accounts” on page 36

Create new or designate existing directory server installation accounts for Process Engine, as shown in the following table.

“Creating Application Engine or Workplace XT accounts” on page 37

Create new or designate existing directory server accounts for Application Engine or Workplace XT.

“Creating Content Search Engine accounts” on page 38

If you are installing Content Search Engine, you must create new or designate existing Autonomy K2 security accounts.

Creating Content Engine directory server accounts:

Create new or designate existing directory server installation accounts for Content Engine.

The Security Administrator should perform this task.

To create directory server accounts:

Create the following users and groups:

| User or group name | Description |
|--|--|
| Bootstrap user name: <i>ce_bootstrap_admin</i> | <p data-bbox="935 394 1417 541">A directory service and application server account that is stored in the <code>CEMPBoot.properties</code> file that is archived in the Content Engine EAR file. Also known as Content Engine System User.</p> <p data-bbox="935 562 1417 709">Content Engine uses <i>ce_bootstrap_admin</i> to establish a connection with the application server, access the application server's JNDI tree, and look up the data sources for accessing the GCD.</p> <p data-bbox="935 730 1417 1024">You should not use this account as an all-purpose administrative account. For example, if you had to log on to some other application using the <i>ce_bootstrap_admin</i> account and provided the wrong password several times, thereby exceeding the number of allowable login failures, this account could be locked out of the directory server, which would mean that Content Engine would not start.</p> <p data-bbox="935 1024 1417 1224">Restriction: If you are deploying Content Engine on an application server with federated user repositories and with multiple realms in your FileNet P8 domain, be sure that no two realms contain the same short name for this user; otherwise, this user will not be able to create the GCD.</p> <p data-bbox="935 1255 1417 1423">See the IBM FileNet P8 help topic System Administration → Enterprise-wide Administration → FileNet P8 Security → Users and groups for information on how <i>ce_bootstrap_admin</i> is the account used to create the initial FileNet P8 domain.</p> |

| User or group name | Description |
|--|--|
| GCD Administrator: <i>gcd_admin</i> | <p>A directory service account that has Full Control access to the domain object for Content Engine.</p> <p>The initial <i>gcd_admin</i> is created by Configuration Manager using the account entered into its Create Bootstrap Properties panel for the Bootstrap user (<i>ce_bootstrap_admin</i>).</p> <p>Log on as <i>gcd_admin</i> in order to:</p> <ul style="list-style-type: none"> • Create the GCD by launching the Configure New Domain Permissions wizard the first time you start Enterprise Manager to establish the IBM FileNet P8 domain (see in the <i>IBM FileNet P8 Platform Installation and Upgrade Guide</i>). • Carry out administrative tasks for the FileNet P8 domain. <p>For more information, see the IBM FileNet P8 help topic System Administration → Enterprise-wide Administration → FileNet P8 Security → Users and groups.</p> |
| Object Store Administrator: <i>object_store_admin, object_store_admin_group</i> | <p>A directory service account that can administer an object store by having Full Control access to it. You can also grant Full Control to an object store to group accounts, thereby making all members of the group object store administrators.</p> <p>Use Enterprise Manager's Create an Object Store wizard to specify which user or group accounts should be <i>object_store_admin_group</i> (see <i>inst_ce_create_obj_store</i> in the <i>IBM FileNet P8 Platform Installation and Upgrade Guide</i>). Each object store could have a different set of object store administrators, depending on your security design.</p> |

| User or group name | Description |
|--|---|
| Directory service user: Sun Java System Directory Server (Referred to as Directory service bind user name in Configuration Manager.): <i>ce_service_user</i> | <p>A Sun Java System Directory Server user account that Content Engine uses to connect to the Sun Java System Directory server. <i>ce_service_user</i> performs the following roles:</p> <ul style="list-style-type: none"> • Acts as the bind user specified by the application server to search through realms to authenticate a user when the user logs in to a Content Engine client such as Workplace. • Acts as the user specified in the GCD that searches users and groups to authorize access to a specific FileNet P8 object after a user has been authenticated. <p>Using Sun Java System Directory Server tools, grant <i>ce_service_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> • Read • Search • Compare <p>Provide the fully qualified distinguished name of this account as the LDAPBindDN while running the Configuration Manager tool and also when you run the Enterprise Manager Directory Configuration Wizard.</p> |

Creating Process Engine directory server accounts:

Create new or designate existing directory server installation accounts for Process Engine, as shown in the following table.

This task should be performed by the Security Administrator. Process Engine can be installed by either a domain or local user. If you choose to install and optionally run Process Engine as a domain user, create the users identified here as optional. You must also create local users and groups for Process Engine installation and the domain users must be added to the local groups. See “Creating Process Engine operating system accounts” on page 13 for procedures for creation of local users and groups.

Create the following users and groups:

| User or Group | Description |
|---|---|
| Process Engine service user: <i>pe_service_user</i> | <p>Process Engine uses the <i>pe_service_user</i> when connecting to the Content Engine server. This user must belong to the Process Engine Administrator group.</p> <p>Record this value in the Installation and Upgrade Worksheet. The name will be required input when configuring the Process Task Manager.</p> |

| User or Group | Description |
|---|--|
| Process Engine administrators group: <i>pe_admin_group</i> | Members of this group automatically have administrative privileges for Process Engine. Record this value in the Installation and Upgrade Worksheet. The name will be required input when configuring the Process Task Manager. |
| Process Engine configuration group: <i>pe_config_group</i> | (Optional) A valid group name. Members of this group automatically have configuration privileges for the Process Engine workflow database. If this group is used to configure security on Process Task Manager, members of this group or of the Process Engine Administrator Group can make configuration changes to the workflow database. If the Process Engine Configuration group is not used during this configuration, anyone can make these changes. Record this value in the Installation and Upgrade Worksheet. The name will be input when configuring the Process Task Manager. |

Creating Application Engine or Workplace XT accounts:

Create new or designate existing directory server accounts for Application Engine or Workplace XT.

This task should be performed by the Application Engine Administrator.

To create Application Engine or Workplace XT accounts:

Create the following users and groups. All IBM FileNet Workplace accounts, as well as accounts for other client applications and expansion products that use Content Engine or Application Engine, must have passwords.

| User or group name | Description |
|---|---|
| Application Engine or Workplace XT installer account (UNIX): <i>ae_install_user</i> , <i>wpxt_install_user</i> | The account you will use to log on to a UNIX machine and launch the Application Engine or Workplace XT installation program. This account must have read/write/execute access to the directory where you will install Application Engine or Workplace XT. |
| Application server account: <i>ae_deploy_user</i> , <i>wpxt_deploy_user</i> | This account will have permissions to deploy an application. The account can be the same as the Application Engine or Workplace XT installer account. |

| User or group name | Description |
|---|---|
| Application Engine Administrators: <i>ae_admin_user, wpxt_admin_user</i> | These accounts will serve in the role of Application Engine administrator. The role applies to both Application Engine or Workplace XT. You will specify these accounts as members of the Application Engine administrator role when you set bootstrap preferences. These accounts must have passwords. |

In addition to the requirements above, the installer account (*ae_install_user* or *wpxt_install_user*) and the Application server account (*ae_deploy_user* or *wpxt_deploy_user*) and need read/write/execute permission to these directories and files:

Installation paths

Grant *ae_install_user* read and write permission to the *ae_install_path*.

Grant *wpxt_install_user* read and write permission to the *wpxt_install_path*.

WebLogic 9.x

BEA_Home/boa/user_projects/domains/domain_name/bin/startWebLogic.sh
or *startWebLogic.cmd*

BEA_home/boa/user_projects/domains/domain_name/config/config.xml

WebLogic 10.x

BEA_Home/boa/wlserver_10.0/server/bin/startWLS.sh or *start WLS.cmd*

BEA_home/boa/user_projects/domains/domain_name/config/config.xml

Creating Content Search Engine accounts:

If you are installing Content Search Engine, you must create new or designate existing Autonomy K2 security accounts.

The IT Administrator should perform this task.

To create accounts for Content Search Engine:

Create the following users and groups:

| User or group name | Description |
|--|---|
| K2 Security Group: <i>k2_sec_group</i> | Autonomy K2 security group used to secure K2 collections. Later, you will specify this group in the User Group field of the Verity Domain Configuration tab when you configure CBR in the Enterprise Manager root domain property sheet. |
| | UNIX UNIX operating system group. |

| User or group name | Description |
|---|--|
| K2 Security User: <i>k2_sec_user</i> | Autonomy K2 security user account, used by Content Engine when logging onto the Autonomy K2 Master Administration Server to perform Content-Based Retrieval (CBR). Later, you will specify this account in the Verity Username field in the Verity Domain Configuration when you configure CBR in Enterprise Manager. This user must be a member of the <i>k2_sec_group</i> and must be defined as an authorized K2 administrator in the K2 dashboard. The <i>k2_sec_user</i> and <i>k2_os_user</i> can be the same user. All permissions listed for both users must be assigned. |
| K2 Operating System User: <i>k2_os_user</i> | <p>UNIX processes will run as this user. <i>k2_os_user</i> can be an unprivileged user; however, the <i>vspget</i> process must run as the root user. (The <i>IBM FileNet P8 Platform Installation and Upgrade Guide</i> explains how to do this.)</p> <p>This user also requires read/write permissions to the K2 collections directory, the collections temp directory and to the file system that contains the file storage areas and full-text index collections. The Autonomy K2 software needs to read the file storage areas and write the full text index collections as part of the full text indexing operation.</p> |

Database administrator installation tasks

The Database administrator must prepare the databases required for IBM FileNet P8, including gathering information about data sources, creating databases and database accounts, and installing client software.

- Review all rows assigned to the Database administrator (DBA) in the Installation and Upgrade Worksheet. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation. (Your organization might have different roles, and some of the responsibilities of listed roles will vary from those assigned by default in this documentation.)

Tip: With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (*p8_worksheet.xls*), perform the following actions to quickly see only the properties assigned to a particular Role:

- Click the **AutoFilter** drop-down arrow in the **Role** column header and select **DBA**.
- Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting **(All)**.

As an alternative, you can use the Customize Worksheet filtering macro, embedded in the worksheet file's Instructions tab.

- If you are installing in a non-English environment, review the considerations and procedures in Appendix A, "Preparing non-English environments for installing IBM FileNet P8 Platform," on page 63 before you begin your preparation tasks.

“Creating Content Engine database accounts”

Use your database tools to create new or designate existing database accounts for Content Engine, as shown in the following table.

“Creating Process Engine database accounts” on page 41

You must create new or designate existing database accounts for Process Engine.

“Preparing Oracle server” on page 42

To prepare your Oracle databases, you must install the database software, and configure database components for IBM FileNet P8.

“Configuring Oracle client for Process Engine” on page 55

Install the Oracle Client software to prepare for the installation of Process Engine. Content Engine does not require installation of Oracle Client software.

“Verifying the ability to connect to the database” on page 56

Verify the ability to connect to the Process Engine database. These procedures can be run after initial configuration of the database and immediately prior to installation of Process Engine software. Run these steps on the database server or the client according to whether the database is local to or remote from Process Engine.

Creating Content Engine database accounts

Use your database tools to create new or designate existing database accounts for Content Engine, as shown in the following table.

Create the following users and groups:

| User or group name | Description |
|--|--|
| Database user name (Oracle): <i>ce_db_user</i> | <p>The tablespace owner accounts that Content Engine uses to access Oracle. Use one account for each object store tablespace and one for the GCD tablespace.</p> <p>Grant each <i>ce_db_user</i> at least the following permissions:</p> <ul style="list-style-type: none"> • CREATE SESSION • CREATE TABLE • SELECT on pending_trans\$ • SELECT on dba_2pc_pending • SELECT on dba_pending_transactions • SELECT on DUAL • SELECT on ALL_TAB_COLS (GCD creation only) • CREATE SEQUENCE (Object Store creation only) • SELECT on product_component_version (Object Store creation only) <p>In addition to the above:</p> <ul style="list-style-type: none"> • On Oracle 10.2.0.4, Oracle 11.1.0.6.0 or later (where both the JDBC client and the database server are at those levels) grant EXECUTE on dbms_xa • On Oracle releases earlier than 10.2.0.4, see Oracle Metalink document ID 436362.1 for required patches, or grant EXECUTE on dbms_system <p>Several of these permissions are required by Content Engine JDBC XA transactions.</p> |

Creating Process Engine database accounts

You must create new or designate existing database accounts for Process Engine.

You create database users on Oracle and SQL Server databases for Process Engine by running SQL scripts. The scripts create default user names and passwords, or they can be edited to change the user names and passwords for your installation. Database permissions are also granted from the SQL scripts.

Create the following users and groups:

| User or group name | Description |
|--|--|
| The database runtime user for Process Engine (Oracle): <i>f_sw</i> | <p>The database runtime user for Process Engine. The user is defined when the <i>pe_filenet_site.sql</i> script runs.</p> <p>The default value is set to <i>f_sw</i>. The script can be edited to change the user name.</p> <p>Record the value for this user in the Installation and Upgrade Worksheet.</p> |

| User or group name | Description |
|---|--|
| The password for the database runtime user for Process Engine (Oracle): <i>f_sw_password</i> | <p>The database runtime user's password for Process Engine. The password is set when the <i>pe_filenet_site.sql</i> and <i>pe_install_scripts.sql</i> script runs.</p> <p>The default value is set to <i>filenet</i>. The script can be edited to change the password.</p> <p>Record the value for this user in the Installation and Upgrade Worksheet.</p> |
| The database maintenance user for Process Engine (Oracle): <i>f_maint</i> | <p>The database maintenance user for Process Engine. The user is defined when the <i>pe_filenet_site.sql</i> script runs.</p> <p>The default value is set to <i>f_maint</i>. The script can be edited to change the user name.</p> <p>Record the value for this user in the Installation and Upgrade Worksheet.</p> |
| The password for the database maintenance user for Process Engine (Oracle): <i>f_maint_password</i> | <p>The database maintenance user's password for Process Engine. The password is set when the <i>pe_filenet_site.sql</i> and <i>pe_install_scripts.sql</i> script runs.</p> <p>The default value is set to <i>change\$this_obnoxio\$ _password</i>. The script can be edited to change the password.</p> <p>Record the value for this user in the Installation and Upgrade Worksheet.</p> |

Preparing Oracle server

To prepare your Oracle databases, you must install the database software, and configure database components for IBM FileNet P8.

“Oracle planning considerations”

Plan the Oracle installation and configuration. Some rules apply to sharing of Oracle databases and tablespaces.

“Verifying that Oracle server is installed for IBM FileNet P8” on page 43

Install the Oracle software and configure the database components for your IBM FileNet P8 system.

“Running Process Engine pre-installation scripts (Oracle)” on page 50

You must run several SQL scripts to create Process Engine users, passwords, tablespace defaults and privileges for the users, stored procedures, and synonyms for the stored procedures. These scripts for Oracle databases can be run in one of three ways.

Oracle planning considerations:

Plan the Oracle installation and configuration. Some rules apply to sharing of Oracle databases and tablespaces.

General

- Determine whether you want to use a dedicated or shared database. In this regard:

- Content Engine, Process Engine, and Rendition Engine can share a database engine, or they can each have a dedicated (unique) database engine.
- Content Engine, Process Engine, and Rendition Engine can each have a dedicated Oracle database instance, or they can share a database instance with one another or with non-IBM FileNet P8 applications.
- Content Engine, Process Engine, and Rendition Engine must each have their own tablespaces, even if they share a database engine.
- Each Process Engine isolated region configured for region recovery must have dedicated tablespaces.

For detailed information regarding installation of Rendition Engine, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

- Plan to use locally managed tablespaces. For performance reasons, IBM recommends that you create locally managed, rather than dictionary managed, tablespaces for Process Engine and Content Engine. (The tablespaces you create via Oracle Database assistant (dbca) are locally managed by default.)
- Be aware that Oracle security defaults have changed between releases. For example, in the Oracle 11g release the default password expiration period is 180 days. Verify the policies are correct for your environment.
- Be aware of database client software requirements. For Process Engine, if the database is remote, you must install database client software on the Process Engine server.

Process Engine

- Process Engine does not support Oracle Password Complexity Verification during the installation process. During installation this Oracle feature must be turned off if the Process Engine run-time user (f_sw or alias) or maintenance user (f_maint or alias) will use default passwords. After installation is complete and the passwords are changed, Oracle's password complexity verification can be turned back on.
- Oracle SQL scripts must be run. A number of SQL scripts must be run that:
 - Create Oracle database accounts for IBM FileNet P8 Process Engine use.
 - Create a number of stored procedures.
 - Grant access levels to the default tablespaces specified in the Process Engine installation program.
- These scripts can be run manually, before starting Process Engine installation, or run from the Process Engine installation program. The Process Engine installation program will complete only if all these scripts run successfully.

Verifying that Oracle server is installed for IBM FileNet P8:

Install the Oracle software and configure the database components for your IBM FileNet P8 system.

You can install and configure an Oracle database that is dedicated or shared by one or more IBM FileNet P8 components. An Oracle database can be:

- dedicated to Content Engine
- dedicated to Process Engine
- dedicated to Rendition Engine
- shared by two or more of Content Engine, Process Engine, and Rendition Engine

In the shared configuration, the IBM FileNet P8 components use the same database, but different table spaces.

You can also share the database with other non-IBM FileNet P8 applications. In the dedicated configuration, Content Engine, Process Engine, and Rendition Engine use separate databases.

A database is local if it is on a machine where you will also be installing Content Engine, Process Engine, or Rendition Engine. A database is remote if it is on a separate server from the component using that database.

For information regarding installation of Oracle Server and Rendition Engine, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

- Make sure the machine that will host the database satisfies all pre-installation requirements specified in the Oracle installation documentation.
 - For Content Engine and Process Engine, IBM FileNet P8 supports the Oracle Advanced Security functionality of secure data transfer across network protocol boundaries.
 - If you will be installing Process Engine on a UNIX machine hosting the Oracle database, be sure that the value of the Oracle environment variable ORACLE_HOME (the path name for the Oracle Server software) is a string of at most 53 characters. If the string has more than 53 characters, the Process Engine installation program will not find the Oracle software, causing the installation to fail.
 - Refer to the *IBM FileNet P8 Hardware and Software Requirements* guide for required operating system and database patch sets, and service packs. The Oracle patches are available at the Oracle Web site. The Oracle patch installation procedure might be less complicated if you do it before you create any databases.
 - Transaction Processing is the required configuration type for the database that supports Content Engine. Choose this configuration type if your database will be dedicated to Content Engine or shared with Process Engine.
 - Record the values for the following settings as you work through the database installation. Enter this information in the Process Engine installer or Content Engine installer sections of the Installation and Upgrade Worksheet, as appropriate. This information must be entered during subsequent installations. Be aware that the Process Engine installation program allows only alphanumeric characters and underscores.
 - Oracle Home
 - Global Database Name
 - Oracle temporary tablespace name
 - Oracle data tablespace name
 - Oracle index tablespace name (optional)
 - Oracle SID
- “Installing an Oracle database engine” on page 45
Install the Oracle software and configure the database server.
- “Configuring Oracle database character sets” on page 46
Determine the Oracle database character set to use for the Content Engine and Process Engine databases.
- “Creating an Oracle database” on page 46
Create one or more databases, depending on whether one or more IBM FileNet P8/IBM FileNet P8 components will share the database.
- “Creating tablespaces for the GCD” on page 46
Create an Oracle tablespaces for the GCD database.

“Creating Oracle tablespaces for Content Engine object stores” on page 47
Create Oracle tablespaces for a Content Engine object store.

“Creating tablespaces for Process Engine” on page 47
Create Oracle tablespaces for Process Engine.

“Setting environment variables for the Oracle and root users on UNIX database server” on page 49

If your Process Engine Oracle database runs on a UNIX machine, set several environment variables in the .profile, .cshrc, or .login file before using the Oracle database. On Windows, the Oracle Universal Installer sets these variables.

“Configuring automatic transaction recovery” on page 49

In a distributed database environment, Oracle MTS Recovery Service (automatically installed with Oracle Services for Microsoft Transaction Server) can resolve in-doubt transactions on the computer that started the failed transaction.

“Turning off Oracle password complexity verification” on page 49

Process Engine does not support Oracle Password Complexity Verification during the installation process if default passwords are used for the database runtime and maintenance users.

“Configuring Oracle XA transactions” on page 50

Configure Oracle XA transactions for Content Engine by running several Oracle SQL scripts.

Installing an Oracle database engine:

Install the Oracle software and configure the database server.

The following procedure shows the minimal choices (specific to the needs of Content Engine and Process Engine) for installing a database engine. Consult Oracle installation documentation for complete preinstallation requirements and instructions. To install an Oracle database engine:

1. Choose the following from the list of available product components:
 - Oracle Server
 - Oracle Net Services
 - Oracle Net Listener
 - Oracle Development Kit
 - Oracle Call Interface (OCI)
 - Oracle Documentation (recommended)
2. If a database is dedicated to Content Engine, or if it will be shared by Content Engine and Process Engine, then Transaction Processing (also known as OLTP) is the required configuration type.
3. If you are going to install Process Engine on this machine, verify/add/edit/uncomment the following lines in the file sqlnet.ora (create the file if it doesn't exist) while the Oracle services/processes are stopped:

```
NAMES.DIRECTORY_PATH=(TNSNAMES)
SQLNET.AUTHENTICATION_SERVICES=(NTS)
```

If Oracle is configured to use LDAP, TNSNAMES must appear in the names.directory_path ahead of LDAP or ONAMES, for example:
NAMES.DIRECTORY_PATH=(TNSNAMES,ONAMES,LDAP).

sqlnet.ora is typically in \$ORACLE_HOME/network/admin on UNIX or ORACLE_HOME\network\admin on Windows operating systems.

4. Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating-system and database patch sets, and service packs.
5. Start the listener and the Oracle database service/processes (Windows/UNIX) if they have not started automatically.

Configuring Oracle database character sets:

Determine the Oracle database character set to use for the Content Engine and Process Engine databases.

- If a database is dedicated to Content Engine or shared with Process Engine, the required character set is AL32UTF8 (Unicode 3.1 UTF-8 universal character set).
- If an object store that is being upgraded uses a database with a character set other than AL32UTF8, you must convert that character set to AL32UTF8. When the character set is not AL32UTF8, applications that attempt to insert Unicode data, for example, when cutting and pasting from a word processing application into Workplace document properties, might (for some characters) experience data loss when stored in the database.
- If the database is dedicated to Process Engine, choose a UNICODE, multibyte, or single-byte 8-bit character set.

Be aware of the special NLS_LANG settings for character sets and locale on Process Engine machines. Consider the following requirements:

- You must ensure that the NLS_LANG Oracle environment variable on an Process Engine machine matches the character set/locale of the operating system.
- Under UNIX, the Oracle Client installer does not automatically set the NLS_LANG. For this reason, after you have installed Process Engine, you must manually set the locale and character set value on each UNIX Oracle Client machine from which users will access Process Engine.
- Set the NLS_LANG value manually on Oracle Client machines using one of the following methods:
 - (UNIX) Add NLS_LANG to the shell environment login files for each user who will be logging on to the machine to run IBM FileNet P8 software.

Creating an Oracle database:

Create one or more databases, depending on whether one or more IBM FileNet P8/IBM FileNet P8 components will share the database.

IBM FileNet P8 requires the following settings for Oracle databases:

Database configuration type

If a database is dedicated to Content Engine, or if it will be shared by Content Engine and Process Engine, then Transaction Processing (also known as OLTP) is the required configuration type.

Server process type

Dedicated Server Mode

Database character set

Choose a database character set, specified in “Configuring Oracle database character sets.”

Creating tablespaces for the GCD:

Create an Oracle tablespaces for the GCD database.

1. Using Oracle Enterprise Manager or SQL*Plus, create a user (*ce_db_user*), password, and default tablespace in the Oracle database for the GCD that Content Engine will access. See Creating Content Engine database accounts for information about the user and required permissions. If one or more object store tablespaces and the GCD tablespace share a physical database, they must not share the same Oracle user. Create different, unique Oracle users for each in this case.
2. Tablespace names must contain only alphanumeric and underscore characters. Names must start with an alphabetic character and must be at most 18 characters long. For performance reasons, you should specify locally managed, instead of dictionary managed, tablespaces. (The tablespaces you create via Oracle Enterprise Manager are locally managed by default.)

The following table shows the recommended minimum sizes of the permanent and temporary tablespaces for each object store that Content Engine will access. (The tablespace names shown in the table are arbitrary.)

| Tablespace Name | Tablespace Type | Minimum Size | Description |
|-----------------|-----------------|--------------|----------------------------------|
| <i>gcd</i> | Permanent | 100 MB | Permanent tablespace for the GCD |
| <i>tempgcd</i> | Temporary | 2 GB | Temporary tablespace for the GCD |

Creating Oracle tablespaces for Content Engine object stores:

Create Oracle tablespaces for a Content Engine object store.

To create Oracle tablespaces:

1. Using Oracle Enterprise Manager or SQL*Plus, create a user (*ce_db_user*), password, and default tablespace in the Oracle database for each object store that Content Engine will access. See Creating Content Engine database accounts for information about the user and required permissions.
2. Create separate users for every object store tablespace.

Tablespace names used by Content Engine must contain only alphanumeric and underscore characters. Names must start with an alphabetic character and must be at most 18 characters long.

For performance reasons, you should specify locally managed, instead of dictionary managed, tablespaces. (The tablespaces you create via Oracle Enterprise Manager are locally managed by default.)

The following table shows the recommended minimum sizes of the permanent and temporary tablespaces for object stores that Content Engine will access. The temporary tablespace can be shared by multiple object stores.

| Tablespace Name | Tablespace Type | Minimum Size | Description |
|------------------|-----------------|--------------|---------------------------------------|
| objectstore1 | Permanent | 200 MB | Permanent tablespace for object store |
| tempobjectstore1 | Temporary | 2 GB | Temporary tablespace for object store |

Creating tablespaces for Process Engine:

Create Oracle tablespaces for Process Engine.

To create tablespaces:

1. Create the tablespaces shown in the following table for the Process Engine. Note that the index tablespace (*vwindex_ts*) is optional. Tablespace names used by Process Engine can contain only alphanumeric and underscore characters. Names must start with an alphabetic character and must be at most 18 characters long.
2. If you will use the region recovery feature in Process Engine, you must create the default tablespaces and an additional data and index tablespace for each region to be configured for recovery. The same runtime and maintenance users will be used for all Process Engine tablespaces.

The following table shows the recommended tablespace names, types, and minimum sizes:

| Tablespace Name | Tablespace Type | Minimum Size (MB) | Description |
|------------------------------|-----------------|-------------------|--|
| <i>vwdata_ts</i> | Permanent | 200 | This is the default name of the data tablespace dedicated to Process Engine. This is the only data tablespace configured during Process Engine installation. Record the tablespace name in the Installation and Upgrade Worksheet. |
| <i>vwtemp_ts</i> | Temporary | 400 | This is the default name of the temporary tablespace dedicated to Process Engine. This is the only temp tablespace configured during Process Engine installation. Record the tablespace name in the Installation and Upgrade Worksheet. |
| <i>vwindex_ts</i> (optional) | Permanent | 200 | This is the default name of the optional default index tablespace. This is the only index tablespace configured during Process Engine installation. If you don't create an index tablespace the data tablespace will be used for indexes. Record the tablespace name in the Installation and Upgrade Worksheet. |

| Tablespace Name | Tablespace Type | Minimum Size (MB) | Description |
|-----------------|-----------------|-------------------|--|
| region X data | Permanent | 200 | This is the data tablespace to be used by an individual region configured for recovery. This tablespace cannot be shared by any other region. |
| region X index | Permanent | 200 | This is the index tablespace to be used by an individual region configured for recovery. This tablespace cannot be shared by any other region. |

Setting environment variables for the Oracle and root users on UNIX database server:

If your Process Engine Oracle database runs on a UNIX machine, set several environment variables in the `.profile`, `.cshrc`, or `.login` file before using the Oracle database. On Windows, the Oracle Universal Installer sets these variables.

To set environment variables:

| Option | Description |
|----------------------------|---|
| For the Oracle user | Set: <ul style="list-style-type: none"> • ORACLE_SID • ORACLE_HOME • Set LD_LIBRARY_PATH to: <code>\$ORACLE_HOME/lib32</code> (Solaris only) |
| For the root user | Set ORACLE_HOME. |

Configuring automatic transaction recovery:

In a distributed database environment, Oracle MTS Recovery Service (automatically installed with Oracle Services for Microsoft Transaction Server) can resolve in-doubt transactions on the computer that started the failed transaction.

- Enable automatic transaction recovery by performing the tasks shown in the section “Scheduling Automatic Microsoft Transaction Server Recovery” in *Oracle Services for Microsoft Transaction Server Developer’s Guide* (Oracle Part Number A95496-01).
- If you are using an Oracle Fail Safe configuration, perform the procedure shown in “Modifying Registry Values for Oracle Fail Safe Configurations” in *Oracle Services for Microsoft Transaction Server Developer’s Guide* (Oracle Part Number A95496-01).

Turning off Oracle password complexity verification:

Process Engine does not support Oracle Password Complexity Verification during the installation process if default passwords are used for the database runtime and maintenance users.

If you are using the default passwords for the runtime (*f_sw*) or maintenance (*f_maint*) users, turn off this Oracle feature and do not re-enable it until you have installed and configured Process Engine.

Configuring Oracle XA transactions:

Configure Oracle XA transactions for Content Engine by running several Oracle SQL scripts.

To configure XA transactions:

1. Log on the Oracle database as either SYSOPER or SYSDBA.
2. Locate and run the `initxa.sql` script in the `ORACLE_HOME\javavm\install` directory.
3. If the script fails to run because the database memory space is too small, locate and run the `initjvm.sql` script in the `ORACLE_HOME\javavm\install` directory. Additional memory-related parameters might need to be adjusted to successfully run this script.

Running Process Engine pre-installation scripts (Oracle):

You must run several SQL scripts to create Process Engine users, passwords, tablespace defaults and privileges for the users, stored procedures, and synonyms for the stored procedures. These scripts for Oracle databases can be run in one of three ways.

- Run them manually on the database server, before running the Process Engine installation program. If the scripts are run manually, both the default runtime and maintenance users and their passwords can be modified in the scripts before execution.
- Run them automatically from the Process Engine installation program, allowing the installation program to prompt for the sys password for Oracle.
- Run them automatically from the Process Engine installation program, running silently using operating system authentication. Use operating system authentication only in a trusted environment or when configured with a local database.

“Scenarios for manually running SQL scripts for a Process Engine Oracle database” on page 51

Several scenarios are used to describe Process Engine installation and configuration parameters and how to set them in each case.

“Process Engine SQL scripts for Oracle” on page 52

All Process Engine SQL scripts are located in the root of the Process Engine installation directory.

“Changes made by `pe_filenet_site.sql` and `pe_ora_users_defaults.sql`” on page 53

The information in columns 1, 2 and 3 is created or set by SQL script `pe_filenet_site.sql`. The information in column 5 is set by SQL script `pe_ora_users_defaults.sql`.

“Changes made by `pe_create_stored_procedures.sql` and `pe_grant_sp_permissions.sql`” on page 53

The information in columns 1, 2 and 3 is created by `pe_create_stored_procedures.sql`. The information in columns 4 and 5 is created by `pe_grant_sp_permissions.sql`.

“Editing SQL scripts and changing default passwords (Oracle)” on page 54
To change the default passwords in an Oracle database environment, you must edit the `pe_filenet_site.sql` and `pe_install_scripts.sql` script files.

“Running SQL scripts manually for a new installation (Oracle)” on page 54
To run SQL scripts for a new installation, copy the scripts to the database server, start SQL Plus, and enter `@pe_install_scripts.sql run_time_user maintenance_user data_tablespace index_tablespace temp_tablespace`.

Scenarios for manually running SQL scripts for a Process Engine Oracle database:

Several scenarios are used to describe Process Engine installation and configuration parameters and how to set them in each case.

Scenario 1: Run SQL scripts manually before installing using all default users and passwords

The DBA wants to run the required SQL scripts manually before running the Process Engine installation program. All default users and passwords are defined in the scripts.

1. Turn off Oracle password complexity.
2. Run the scripts manually, before running the Process Engine installation program, without making any changes in the scripts to passwords, setting the runtime user to `f_sw` and the maintenance user to `f_maint`.
3. Run Process Engine installation program and use default users for the database accounts (do not set aliases).
4. Leave the `f_sw` and `f_maint` password fields blank (using the defaults assigned with the scripts ran).
5. Turn Oracle password complexity back on.
6. Reset the `f_sw` and `f_maint` passwords by running `Xdbconnect`. This changes both the encrypted version of the password and the password in the database. `Xdbconnect` works only if the passwords in the encrypted file and the database match after installation.

Scenario 2: Run SQL scripts manually before installing with non-default users and passwords

The DBA wants to run the required SQL scripts manually before running the Process Engine installation program, but non-default users and passwords are set.

1. Do not turn off Oracle password complexity.
2. Edit the scripts to change the passwords.
3. Run the scripts, entering the desired runtime and maintenance users and passwords.
4. Run the Process Engine installation program and indicate that aliases will be configured. The user names set when the scripts ran must be defined as the aliases for `f_sw` and `f_maint`. None of the fields for alias names can be left blank in the Process Engine installation screen, but default values can be entered where default values are still wanted.
5. Set the `f_sw` and `f_maint` passwords during Process Engine installation to the match the passwords set when the scripts were run manually.
6. Do not reset passwords by running `Xdbconnect`. Because non-default passwords were used, there is no need to change them immediately after installation.

Process Engine SQL scripts for Oracle:

All Process Engine SQL scripts are located in the root of the Process Engine installation directory.

| Script Name | Action |
|---------------------------------|---|
| pe_install_scripts.sql | <p>A wrapper script that runs the following scripts: pe_filenet_site.sql, pe_create_stored_procedures.sql, and pe_grant_sp_permissions.sql.</p> <p>The pe_install_scripts.sql script generates an Oracle output spool file. After the script completes, the Process Engine installation program scans the output file for errors.</p> |
| pe_filenet_site.sql | <p>Creates the database users for IBM FileNet P8. See “Changes made by pe_filenet_site.sql and pe_ora_users_defaults.sql” on page 53 for details.</p> |
| pe_create_stored_procedures.sql | <p>Creates several stored procedures. See “Changes made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql” on page 53 for details.</p> |
| pe_grant_sp_permissions.sql | <p>Runs grants and creates synonyms for stored procedures. See “Changes made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql” on page 53 for details.</p> |
| pe_ora_users_defaults.sql | <p>Sets user tablespace defaults and privileges for IBM FileNet P8 users. See “Changes made by pe_filenet_site.sql and pe_ora_users_defaults.sql” on page 53 for details.</p> |

The Process Engine scripts run with the following options:

run_time_user

This option value is one of the following:

- f_sw - if you use the default database users when you run the Process Engine installation program.
- alias for f_sw - if you define aliases for the operating system and database users when you run the Process Engine installation program.

maintenance_user

This option value is one of the following:

- f_maint - if you use the default database users when you run the Process Engine installation program.
- alias for f_maint - if you define aliases for the operating system and database users when you run the Process Engine installation program.

data_tablespace (Oracle only)

The data tablespace name that will be entered during Process Engine installation.

index_tablespace (Oracle only)

The index tablespace name that will be entered during Process Engine installation. This is an optional tablespace. If it does not exist, enter the data tablespace value.

temp_tablespace (Oracle only)

The temp tablespace name that will be entered during Process Engine installation.

Changes made by pe_filenet_site.sql and pe_ora_users_defaults.sql:

The information in columns 1, 2 and 3 is created or set by SQL script pe_filenet_site.sql. The information in column 5 is set by SQL script pe_ora_users_defaults.sql.

| FileNet DB user created | Granted permissions/privileges/roles | Default Initial password | Can delete post-install? | Default, temp, and index Tablespace privileges*** |
|-------------------------|---|---------------------------------|--------------------------|---|
| f_sw or alias | Create session, alter session, create table, create view, create sequence, create public synonym, drop public synonym, create procedure Select on sys.dba_users Select on sys.dba_tablespaces Create public synonym Drop public synonym | filenet | No | Data tablespace set Temp tablespace set Index tablespace set Quota 0 on system Quota unlimited on default tablespace Quota unlimited on temp tablespace Quota unlimited on index tablespace |
| f_maint or alias | DBA role | change\$this_obnoxious_password | Yes | |

Passwords can be manually changed after installation with the **Xdbconnect** tool.

*** Default data, temp, and index tablespace names are set during installation. These names are used to set database user privileges. The index tablespace is optional. If not specified, the data tablespace will be used.

Changes made by pe_create_stored_procedures.sql and pe_grant_sp_permissions.sql:

The information in columns 1, 2 and 3 is created by pe_create_stored_procedures.sql. The information in columns 4 and 5 is created by pe_grant_sp_permissions.sql.

| Procedure name | Owner | Description | Grants | Synonym |
|------------------|---------------|---|-------------------|------------------|
| fn_error | f_sw or alias | Displays text of a specified ORA error number. Calls SQL stored procedure fn_errortxt to get the text. | Execute to public | fn_error |
| fn_oraversion | f_sw or alias | Displays version number of Oracle RDBMS. Calls SQL stored procedure fn_oraversiontxt. | Execute to public | fn_oraversion |
| fn_errortxt | f_sw or alias | Gets and returns the message text of a specified ORA SQL Error Code. | Execute to public | fn_errortxt |
| fn_oraversiontxt | f_sw or alias | Gets and returns the version number of the Oracle RDBMS. | Execute to public | fn_oraversiontxt |

Editing SQL scripts and changing default passwords (Oracle):

To change the default passwords in an Oracle database environment, you must edit the `pe_filenet_site.sql` and `pe_install_scripts.sql` script files.

1. Open the `pe_filenet_site.sql` script with a text editor.
2. Locate and edit the default runtime password (filenet) in the following command:

```
grant create session, alter session, create table,
create view, create sequence, create public synonym, drop public
synonym, create procedure to &1 identified by filenet;
```

3. Locate and edit the default maintenance password (change\$this_obnoxio\$_passwrđ) in the following command:
4. Open the `pe_install_scripts.sql` script with a text editor.
5. Locate and edit the default runtime password (filenet) in the following command:

```
connect &1/filenet
```

Running SQL scripts manually for a new installation (Oracle):

To run SQL scripts for a new installation, copy the scripts to the database server, start SQL Plus, and enter `@pe_install_scripts.sql run_time_user maintenance_user data_tablespace index_tablespace temp_tablespace`.

1. From the Process Engine software package, copy the scripts to the database server.
2. Start SQL Plus. For example, enter the command: `sqlplus "sys/password as sysdba"`
3. At the SQL prompt, enter: `@pe_install_scripts.sql run_time_user maintenance_user data_tablespace index_tablespace temp_tablespace` For example:

```
@pe_install_scripts.sql f_sw
f_maint vwdata_ts vwindex_ts vwtemp_ts
```

Configuring Oracle client for Process Engine

Install the Oracle Client software to prepare for the installation of Process Engine. Content Engine does not require installation of Oracle Client software.

For information regarding installation of Oracle Client and Rendition Engine, see the *IBM FileNet Rendition Engine Installation and Upgrade Guide*.

- Install Oracle Client on any machine that will host Process Engine, Rendition Engine, or any other IBM FileNet P8 component except Content Engine (such as Enterprise Manager) that needs to access an Oracle database.
- Install a 32-bit Oracle Client for Process Engine.
- If you will be installing Process Engine on a UNIX machine hosting Oracle Client software, be sure that the value of the Oracle environment variable ORACLE_HOME (the path name for the Oracle Client software) is a string of at most 53 characters. If the string has more than 53 characters, the Process Engine installer will not find the Oracle software, causing the installation to fail.
- Refer to the *IBM FileNet P8 Hardware and Software Requirements* for required operating system and database patch sets, and Service Packs. The Oracle patches are available at the Oracle website. The Oracle patch installation procedure might be less complicated if you do it before you create any databases.

“Installing the Oracle client”

You have minimal choices (specific to the needs of Process Engine) for installing a database client. Consult the Oracle installation documentation for complete preinstallation requirements and instructions.

Installing the Oracle client:

You have minimal choices (specific to the needs of Process Engine) for installing a database client. Consult the Oracle installation documentation for complete preinstallation requirements and instructions.

To install the Oracle client:

1. Choose the following from the list of available product components:
 - Oracle10g Client or Oracle11g Client as appropriate
 - Oracle Network Utilities
 - Oracle Database Utilities
 - SQL*PLUS
2. Using Oracle Net Configuration Assistant, test the connection to the Oracle database server with an appropriate Oracle user and password.
3. If you are going to install Process Engine software on this machine, and your remote Oracle database uses the Unicode character set AL32UTF8, then for each user who will access Process Engine software on the machine, set the value of the Oracle environment variable parameter NLS_LANG to reflect the PE-supported locale and (non-Unicode) character set on the machine.

| Option | Description |
|--------|--|
| UNIX | Add NLS_LANG to the shell environment login files for each user who will be logging onto the machine to run IBM FileNet P8 software. |

To affect the environment for only Process Engine, set NLS_LANG for just the fns user. (On Windows platforms, fns is created by the Process Engine

installer; on UNIX platforms, you manually create fnswh as part of the Process Engine installation task.)

- If you are going to install Process Engine software on this machine to connect to a remote Oracle database, set the value of an environment variable for the oracle user to a default connect identifier, such as the Oracle net service name or the database service name. The name of the environment variable depends on which operating system is on this machine:

| Option | Description |
|--------|-------------|
| UNIX | TWO_TASK |

- If you are going to install Process Engine software on a UNIX machine to connect to a remote Oracle database, set the following environment variables in the startup file of the default oracle user.
 - ORACLE_HOME
 - Set PATH to: \$ORACLE_HOME/bin (UNIX only)

- If you are going to install Process Engine on this machine, verify that the sqlnet.ora file exists and that the following lines are in it.

```
NAMES.DIRECTORY_PATH=(TNSNAMES)
SQLNET.AUTHENTICATION_SERVICES=(NTS)
```

If Oracle is configured to use LDAP, TNSNAMES must appear in the names.directory_path ahead of LDAP or ONAMES, for example:

```
NAMES.DIRECTORY_PATH=(TNSNAMES,ONAMES,LDAP)
```

sqlnet.ora is typically in \$ORACLE_HOME/network/admin on UNIX or ORACLE_HOME\network\admin on Windows operating systems.

- Install all required Oracle patches, as specified in the *IBM FileNet P8 Hardware and Software Requirements*. These patches are available at the Oracle Web site.

Verifying the ability to connect to the database

Verify the ability to connect to the Process Engine database. These procedures can be run after initial configuration of the database and immediately prior to installation of Process Engine software. Run these steps on the database server or the client according to whether the database is local to or remote from Process Engine.

“Verifying the Process Engine database connection (Oracle)”

Verify that the Oracle database instance used by Process Engine is accessible.

How you log on to sqlplus will vary, depending upon how you will choose to run the SQL scripts.

Verifying the Process Engine database connection (Oracle):

Verify that the Oracle database instance used by Process Engine is accessible. How you log on to sqlplus will vary, depending upon how you will choose to run the SQL scripts.

This procedure will verify that you can connect to the Oracle database in the same way the Process Engine installation program will. Make whatever corrections are necessary before proceeding.

To verify the Process Engine database connection:

- Run the following at a command prompt on the Process Engine:

```
su - oracle -c "sqlplus"
```


2. Enter one of the following commands at the SQLPlus prompt, as follows:
 - If the Process Engine pre-installation SQL scripts will be run from the Process Engine installation program by prompting for the sys password, type the following command:
`sys/password as sysdba`
 - If the SQL scripts will be run from the Process Engine installation program by using operating system authentication, type the following command:
`/ as sysdba`
3. At the prompt, enter the following SQL command:
`SQL> select instance_name, host_name, version
from v$instance;`

The following represents an example of the information returned:

```

INSTANCE_NAME
-----
HOST_NAME
-----
VERSION
-----
p8dbshr HQVWBUCS 10.2.0.2.0
where:
  • p8dbshr is the instance ORACLE_SID.
  • HQVWBUCS is the database server name.
  • 10.2.0.2.0 is the Oracle server version.

```

Application Server administrator installation tasks

The Application Server Administrator must prepare the application servers for IBM FileNet P8 Platform, including planning deployment, creating administrative accounts, and configuring JDBC drivers for both Content Engine and Application Engine.

- Review all rows assigned to the Application Server Administrator (ASA) in the Installation and Upgrade Worksheet. While you complete the following preparation tasks, provide values for the rows that are appropriate to your installation.

Tip: With the **Data** → **Filter** → **AutoFilter** command enabled, as it is by default in the shipping worksheet file (p8_worksheet.xls), perform the following actions to quickly see only the properties assigned to a particular Role:

- Click the **AutoFilter** drop-down arrow in the **Role** column header and select ASA.
- Further filter the result set by clicking the **AutoFilter** drop-down arrow in any of the other columns and selecting a value or clear a filter by selecting (All).
- If you are installing in a non-English environment, review the considerations and procedures in Appendix A, “Preparing non-English environments for installing IBM FileNet P8 Platform,” on page 63 before you begin your preparation tasks.
- Review the following planning considerations:
 - “Application server planning considerations” on page 4
 - “Creating Content Engine application server accounts” on page 58
 Create new or designate existing application server accounts for Content Engine, as shown in the following table.

“Configuring WebLogic Server for Content Engine”

You need to configure WebLogic Server after installing it on the machine where you are going to install and deploy Content Engine.

“Configuring WebLogic Server for Application Engine or Workplace XT” on page 60

You must install Oracle WebLogic Server on the machine where you are going to install and deploy Application Engine or Workplace XT.

“Configuring Process Engine clients for ORB” on page 61

Process Engine clients require either the IBM or the Sun Object Request Broker (ORB).

“Configuring the documentation server” on page 61

If you want to have a local documentation installation instead of linking to the information center at www.ibm.com, plan to install IBM FileNet P8 documentation on an application server. Deploying the help as a Web application ensures that you can access online help from within IBM FileNet P8 applications and use the full-text search feature.

Creating Content Engine application server accounts

Create new or designate existing application server accounts for Content Engine, as shown in the following table.

| User/Group | Description |
|--|--|
| Application server administrator user name <i>ce_appserver_admin</i> | Create an application server administrative account that can log on to Content Engine's application server administration console. |
| Administrative Console user <i>ce_appserver_console_admin</i> | |

Configuring WebLogic Server for Content Engine

You need to configure WebLogic Server after installing it on the machine where you are going to install and deploy Content Engine.

Before installing and deploying Content Engine on a WebLogic Server machine, you must create a WebLogic Server domain and install JDBC drivers (The drivers must be installed on the WebLogic Server machine whether your database is collocated or not).

To configure WebLogic Server:

1. Use the WebLogic Configuration Wizard to create a WebLogic Server domain for Content Engine. In the examples below, use the domain name `FNCEDomain`. Keep the following in mind as you configure the domain:
 - a. Set the server start mode to Production mode.
 - b. Select the appropriate Java™ Development Kit (JDK) for your environment, as specified in the *IBM FileNet P8 Hardware and Software Requirements*.
2. Use the WebLogic Administration Console to set the following:
 - a. Optional: Create a WebLogic Server authentication provider. You can use the Content Engine Configuration Manager tool to create a WebLogic Server authentication provider later on, or you can create the provider now by using the WebLogic Console.

Important: In some situations (for example, if you have a single-sign-on provider, such as Netegrity SiteMinder), Configuration Manager cannot configure a WebLogic authentication provider.

- b. For performance reasons set the parameters that control searches within the authentication provider, as shown in the following table:

| Parameter | Value | Description |
|-----------------------------------|-----------|---------------------------------------|
| Group Membership Searching | unlimited | Group searches are unlimited in depth |
| Max Group Membership Search Level | 0 | Only direct group members are found |

If you encounter performance problems, change the Group Membership Searching parameter value to limited.

- c. If you do not use Configuration Manager to configure your LDAP settings, set the Control Flag value. Set the Control Flag to REQUIRED to allow logons to FNCEDomain by LDAP-authenticated users in the Default Authenticator who are not in FNCEDomain's active security realm. The Configuration Manager tool will set this flag when you run the Configure LDAP task.
- d. If you are using multiple authentication providers in an Active Directory environment of multi-forest domains, reorder (as needed) the list of providers so that the most-frequently-used provider is first in the list, and the least-frequently-used is last. Reordering is necessary to prevent logon failures when IBM FileNet P8 Workplace is being accessed by many users simultaneously.
- e. Specify the following heap sizes for the JVM:
- Initial Java heap size (-Xms):**
512 MB
- Maximum Java heap size (-Xmx):**
1024 MB
3. Refer to the *IBM FileNet P8 Hardware and Software Requirements* guide for information on the JDBC driver file for the database type that you need for the GCD or for an object store you will be creating later.
4. Depending on your database, use one of the following procedures to install the JDBC drivers.

Oracle

- a. Check to see if the Oracle JDBC Driver file is already on your WebLogic machine by searching for `ojdbc##.jar` in the `wls_install_path/server/lib` directory, where `wls_install_path` is the WebLogic Server installation path, such as `C:\bea\weblogic92`.
- b. If no Oracle JDBC Driver file is present, download the file (the one that matches the version of the JDK on your WebLogic Server machine) from the Oracle JDBC Driver Downloads web site to a directory on the WebLogic machine.

Restriction: If you intend to install AddOns (extensions to IBM FileNet P8 core components), and your Content Engine database is Oracle, your Oracle JDBC Driver file requirements might be more restrictive. For the required version and patch number, see the *IBM FileNet P8 Hardware and Software Requirements*.

- c. From the Oracle web site, apply the patch Oracle Patch `Ojdbc##.jar`.

- d. Edit the file `startWebLogic.cmd` or `startWebLogic.sh` for the WebLogic domain you created. Add the following lines immediately after the first line that starts with `set CLASSPATH`.

UNIX

```
JDBC_PATH=<jdbc_path>/ojdbc##.jar  
CLASSPATH=$JDBC_PATH:$CLASSPATH
```

5. If your application server uses the IBM JVM, edit the `JAVA_OPTIONS` variable to improve performance.

UNIX Immediately before the `SAVE_JAVA_OPTIONS=$JAVA_OPTIONS` line in the `startWebLogic.sh` file, insert the following line. Do not type any line breaks.

```
JAVA_OPTIONS="$JAVA_OPTIONS -Dprogram.name=$PROGNAME  
-Dfilenet.pe.peorb.pool.min=2 -Dfilenet.pe.peorb.pool.max=5"
```

6. Stop and then start WebLogic Server.
7. Give the Configuration Manager user, *config_mgr_user*, the following permissions:
 - a. Read, write, and execute permission on the domain directory `./users_projects/domains/your_domain`.
 - b. Read and execute permission on the `./common/bin` directory.

Configuring WebLogic Server for Application Engine or Workplace XT

You must install Oracle WebLogic Server on the machine where you are going to install and deploy Application Engine or Workplace XT.

Application Engine or Workplace XT can be collocated with Content Engine as long as the server is appropriately sized. However, each instance of Application Engine or Workplace XT and each instance of the Content Engine must run in its own JVM. For assistance in sizing your system, contact your service representative.

To configure WebLogic Server:

1. Verify that the application server is set to use `JSESSIONID` as the default cookie name. To avoid forcing end users to log in individually to applets such as Process Designer, Search Designer, and Process Simulator, configure the application server to use `JSESSIONID` as cookie name, and not use application-unique cookie names. Using `JSESSIONID` is typically the default setting for the supported application servers. Both Application Engine and Workplace XT use cookie names for passing session information between Application Engine or Workplace XT and the client browser.
2. Create a WebLogic Server domain before installing and deploying Application Engine or Workplace XT. Refer to your WebLogic Server documentation for detailed instructions.

Important: You will perform further configuration on Workplace XT after you perform the installation.

3. Enable trust between the Content Engine WebLogic domain and the Application Engine or Workplace XT WebLogic domain. If you have already configured a trust on the Content Engine, perform these steps only on the Application Engine or Workplace XT server. You must use the same password as the one used on the Content Engine.
 - a. Prepare the WebLogic Administration Console for editing.

- b. In the advanced security settings for your domain, enter a password for the domain. You must enter the same password for both the Content Engine domain and the Application Engine or Workplace XT domain.
- c. Save and activate your changes.
- d. Restart the server if needed.
- e. Repeat this procedure in each domain for which you want to enable trust.

Important: If you are enabling this feature in a managed server environment, you must stop the Administration server and all the Managed Servers in both domains and then restart them. If this step is not performed, servers that were not rebooted will not trust the servers that were rebooted. Refer to your BEA documentation for more information.

Configuring Process Engine clients for ORB

Process Engine clients require either the IBM or the Sun Object Request Broker (ORB).

This applies to the following configurations:

- J2EE application server clients such as Workplace or Workplace XT
- Content Engine when using the workflow subscription processor to launch workflows
- Non-J2EE and custom applications

The default ORB varies by application server, so in most instances no changes are required. However, in certain configurations you must override the defaults as follows:

| Application server | Changes |
|---|--|
| Oracle WebLogic Server with the Sun JVM | No changes are required. |
| WebLogic Server with the JRockit JVM | Replace the default ORB with the Sun ORB by adding the following code to your application server startup file. Enter without carriage returns: <pre>set JAVA_OPTIONS=%JAVA_OPTIONS% -Dorg.omg.CORBA.ORBClass =com.sun.corba.se.impl.orb.ORBImpl set JAVA_OPTIONS=%JAVA_OPTIONS% -Dorg.omg.CORBA.ORBSingletonClass =com.sun.corba.se.impl.orb.ORBSingleton</pre> |

Configuring the documentation server

If you want to have a local documentation installation instead of linking to the information center at www.ibm.com, plan to install IBM FileNet P8 documentation on an application server. Deploying the help as a Web application ensures that you can access online help from within IBM FileNet P8 applications and use the full-text search feature.

Ensure that you have a supported application server. The application server can be one of the servers that you already prepared for Content Engine or Application Engine.

After you have chosen the application server for the documentation, you can specify the value for the documentation URL. This is useful when recording the installation properties for the other IBM FileNet P8 components.

To determine the documentation URL:

1. Choose the application server for the IBM FileNet P8 documentation.
2. Make a note of the server URL. The format of the URL depends on the documentation package you choose:

Table 1. Documentation URLs

| FileNet P8 Documentation package | URL |
|---|--|
| Installed Eclipse-based information center | <code>http://yourdocserver:port#/p8docs/index.jsp</code> |
| Installed non-Eclipse-based documentation package | <code>http://yourdocserver:port#/ecm_help/_start_here.htm</code> |

3. Search the Installation and Upgrade Worksheet for documentation server properties in the **Property** or **Parameter** column.
4. Enter the URL value into the worksheet for each instance of that property for the components you are going to install.

Appendix A. Preparing non-English environments for installing IBM FileNet P8 Platform

To run IBM FileNet P8 components in a non-English environment, certain conditions must be met. Review the following considerations and tasks, organized by administrator role, if you plan to run IBM FileNet P8 in a non-English environment.

“Application Server administrator”

To support Unicode UTF-8 characters, all FileNet P8 domain application servers must be properly configured and must have all fix packs installed.

“Security administrator” on page 64

The IBM FileNet P8 security administrator installation role includes configuring and maintaining directory servers.

“Database administrator” on page 64

The IBM FileNet P8 database administrator installation role includes configuring database installations and tablespaces, and creating database accounts.

“IT administrator” on page 67

Depending on the operating system, the IT administrator installs either a localized version of the operating system, or the operating system language pack.

“IBM FileNet P8 administrator” on page 69

The FileNet P8 administrator starts Process Engine services and configures Process Task Manager for Application Engine and Workplace XT.

“Limitations on installing in a non-English environment” on page 69

There are certain limitations on installing IBM FileNet P8 in non-English environments.

Application Server administrator

To support Unicode UTF-8 characters, all FileNet P8 domain application servers must be properly configured and must have all fix packs installed.

“Configuring character encoding (WebLogic Server)”

Set the appropriate codesets to UTF-8.

“WebLogic Server” on page 64

(WebLogic Server 9.x) To fix the BEA CR298435 character corruption problem in FileNet P8, apply WebLogic Server 9.2 MP1 on WebLogic Server 9.x.

Configuring character encoding (WebLogic Server)

Set the appropriate codesets to UTF-8.

To configure the correct character encoding:

Set the following Interop-Orb-Protocol (IIOP) attributes:

- Default Char Codeset to "UTF-8"
- Default Wide Char Codeset to "UTF-8"

WebLogic Server

(WebLogic Server 9.x) To fix the BEA CR298435 character corruption problem in FileNet P8, apply WebLogic Server 9.2 MP1 on WebLogic Server 9.x.

Security administrator

The IBM FileNet P8 security administrator installation role includes configuring and maintaining directory servers.

“Extended characters and user names”

Note the following considerations for localized IBM FileNet P8 accounts.

Extended characters and user names

Note the following considerations for localized IBM FileNet P8 accounts.

- Process Engine does not support extended characters in LDAP attributes for authentication purposes. These attributes include, but are not limited to, such items as cn (common name), ou (organizational unit), or dc (domain component). ASCII characters are required for these attributes.
 - Process Engine does not support non-ASCII user names. Therefore, do not use non-ASCII user names when starting the Process Task Manager.
 - WebDAV and the SSO environment also do not support Non-ASCII user names.
 - The Content Engine locale must match directory server locale to manage non-ASCII user names correctly.
 - UNIX systems can support Latin1, Latin2, Arabic, and double-byte user names simultaneously.
-

Database administrator

The IBM FileNet P8 database administrator installation role includes configuring database installations and tablespaces, and creating database accounts.

“Installing Oracle server”

Before installing Oracle, verify that the operating system locale is set appropriately.

Installing Oracle server

Before installing Oracle, verify that the operating system locale is set appropriately.

Process Engine

A single instance of Process Engine on an Oracle server in a UNIX UTF-8 locale environment can support multiple languages by selecting the Unicode AL32UTF8 database character set when configuring the Oracle database. If support for more than one language is not required, select the appropriate language database character set instead.

Some examples of Process Engine supported database character sets are:

AL32UTF8

Supports all languages that Process Engine currently supports.

JA16SJIS

Supports only Japanese and English.

WE8ISO8859P15

Supports West European with Euro currency support, but not Japanese.

Restriction: Process Engine does not support the National Character Set AL32UTF16 (NCHAR, NVARCHAR2, and NCLOB columns).

“Installing Oracle client”

Consider the following language information when installing Oracle Client software.

“UNIX” on page 66

For all Process Engine deployments using a remote Oracle database, confirm the following two environment variables in the fns user profile are set to the proper locale for Process Engine.

Installing Oracle client

Consider the following language information when installing Oracle Client software.

“NLS_LANG environment variable - charset”

When data is exported from an Oracle database, it is converted from the database character set to the character set specified in the NLS_LANG environment variable on the Oracle client.

NLS_LANG environment variable - charset:

When data is exported from an Oracle database, it is converted from the database character set to the character set specified in the NLS_LANG environment variable on the Oracle client.

When data is imported to an Oracle database, it is assumed imported data is already in the character set specified in the NLS_LANG environment variable.

If the database character set matches the operating system locale, you do not need to set the NLS_LANG environment variable. However, if the database character set does not match the operating system locale and the NLS_LANG environment variable is not set, data could become garbled because the database will incorrectly map the import/export data between the database and operating system.

Requirement: The NLS_LANG environment variable setting must match the UNIX locale or the Windows code page of the Process Engine server if the Oracle database is remote.

The format of the NLS_LANG environment variable is:

NLS_LANG = language_territory.charset

The following table provides a brief explanation about the environment variable arguments. For more information, see the Oracle documentation.

| Arguments | Description |
|-----------|---|
| Language | Specifies the language used for Oracle messages, day names, and month names. Each supported language has a unique name, for example, American, French, or Japanese. The language argument has default values for the territory and character set arguments, so either (or both) of the territory or the character set arguments can be omitted. If a language is not specified, the value defaults to American. |

| Arguments | Description |
|-----------|---|
| Territory | <p>Specifies the default date format and decimal character used for numbers. Each supported territory has a unique name, for example, America, France and Japan.</p> <p>Because each language provides a default territory, a value is not required.</p> |
| Charset | <p>Specifies the character set used by the client application (normally that of the user's terminal). Each supported character set has a unique acronym, for example, US7ASCII, WE8ISO8859P1, or JA16EUC.</p> <p>Because each language provides a default character set, a value is not required.</p> |

UNIX

For all Process Engine deployments using a remote Oracle database, confirm the following two environment variables in the fns user profile are set to the proper locale for Process Engine.

The LANG environment variable specifies internationalization information that allows users to work with different national conventions in UNIX. For more information, see the UNIX documentation.

“Setting the NLS_LANG environment variables”

The task of setting the NLS_LANG environment variables consists of configuring the language, territory and character set that correspond to the locale of your Process Engine server.

Setting the NLS_LANG environment variables:

The task of setting the NLS_LANG environment variables consists of configuring the language, territory and character set that correspond to the locale of your Process Engine server.

Modify the fns user.profile file in /home/fns:

```
export NLS_LANG=language_territory.charset
```

where: language_territory.charset is the language, territory, and charset that best match your Process Engine server locale.

Example of Process Engine LANG, Oracle character sets, and NLS_LANG configurations:

The following table includes some sample NLS_LANG settings for different Process Engine configurations. Most important is to match your NLS_LANG to a Process Engine locale, not the Oracle database character set. For multilingual Process Engine configurations, the database character set must be AL32UTF8, and the server must use an UTF-8 locale. For example, LANG=JA_JP.UTF-8, or NLS_LANG=Japanese_Japan.UTF8 or NLS_LANG=French_France.UTF8@euro.

| OS | Locale/Windows Code Page | Remote Oracle Character Set | NLS_LANG for PE Server fnsw User | Comments |
|-------------|--------------------------|-----------------------------|----------------------------------|--|
| Sun Solaris | LANG=ja_JP.eucJP | JA16SJIS | NLS_LANG=Japanese_Japan.JA16EUC | Because Solaris data will be in eucJP, the Oracle JA16EUC character set is the best match. |

IT administrator

Depending on the operating system, the IT administrator installs either a localized version of the operating system, or the operating system language pack.

“Operating system considerations”

In addition to any operating system platforms, the IT administrator must consider the FileNet P8 Platform components that will be installed in a non-English environment.

“UNIX” on page 68

Install the operating system language packs. For more information, see the operating system installation documentation.

Operating system considerations

In addition to any operating system platforms, the IT administrator must consider the FileNet P8 Platform components that will be installed in a non-English environment.

Application Engine or Workplace XT

Application Engine and Workplace XT can be installed:

- In any locale on any of the supported UNIX platforms
- On any localized version of Windows or in any region on the English version of Windows

Remember: The Application Engine and Workplace XT setting must match the Process Engine setting when the Process Task Manager is started. Otherwise, workflows can experience unexpected problems such as errors related to the way characters display.

Content Engine

Content Engine can be installed:

- In any locale on any of the supported UNIX platforms
- On any localized version of Windows or in any region on the English version of Windows

Process Engine

Process Engine must be installed on:

- A UTF-8 locale and a UTF-8 database character set when used with Oracle or DB2

- On any localized version of Windows for the supported language or in any region on the English version of Windows
- In UNIX environments, Process Engine language support is determined by a combination of the operating system locale and the database character set.
- The Java Runtime Environment version must be 1.5.0 or higher for the language capability to function correctly.

UNIX

Install the operating system language packs. For more information, see the operating system installation documentation.

“Configuring support for other languages in X Windows UNIX system”

Add language fonts for your UNIX operating system if necessary to display an X-Windows desktop in a specific-language user interface. Follow your UNIX operating system administration guide to install other language fonts.

“Setting the LANG, LC_TIME, and LC_MESSAGES environment variables”

When the IT administrator sets the LANG, LC_TIME, and LC_MESSAGES environment variables, the locale must be the same as the UNIX locale and the Process Engine locale.

Configuring support for other languages in X Windows UNIX system

Add language fonts for your UNIX operating system if necessary to display an X-Windows desktop in a specific-language user interface. Follow your UNIX operating system administration guide to install other language fonts.

Configure your X-session manager application to use UNIX operating system fonts. Refer to your X-session manager application administration guide for details on adding fonts or accessing them on the UNIX server.

Setting the LANG, LC_TIME, and LC_MESSAGES environment variables

When the IT administrator sets the LANG, LC_TIME, and LC_MESSAGES environment variables, the locale must be the same as the UNIX locale and the Process Engine locale.

Set the LANG, LC_TIME, and LC_MESSAGES environment variables:

For example:

```
LANG=locale
export LANG
```

Where *locale* is the same as the UNIX locale and the Process Engine locale.

- Use UTF-8 if Process Engine needs to support multiple languages that are not in the same language code page.
- Use the UNIX "locale -a" command to determine the available locales.
- Setting the LANG variable might reset the LC_TIME and LC_MESSAGES variables to the locale set in LANG. If this happens, you must change the settings of LC_TIME and LC_MESSAGES back to "C" or equivalent ISO-8859 locales.

For example:

```
export LC_TIME="C"
export LC_MESSAGES="C"
```

IBM FileNet P8 administrator

The FileNet P8 administrator starts Process Engine services and configures Process Task Manager for Application Engine and Workplace XT.

“Process Engine services”

All Process Engine services must start in a Unicode locale to support multiple languages on a UNIX platform. Regions initialized in a specific locale must restart in the same locale.

“Configuring Process Task Manager for Application Engine and Workplace XT”

In UNIX environments, verify the operating system locale is same as the Process Engine locale, and verify the LC_TIME and LC_MESSAGES are set to “C” or an equivalent ISO-8859 locale before running Process Task Manager in Application Engine or Workplace XT. Failing to do so can result in character corruption and application failure.

Process Engine services

All Process Engine services must start in a Unicode locale to support multiple languages on a UNIX platform. Regions initialized in a specific locale must restart in the same locale.

Attention: Starting Process Engine services in a different locale can corrupt transferred workflow and region queue names in the isolated region. If an isolated region's non-English queue names are corrupted due to starting Process Engine services in the wrong locale, the only recovery option available is to re-initialize the region to clean up the Process Engine database. Transferred workflows and region queues will be lost.

Configuring Process Task Manager for Application Engine and Workplace XT

In UNIX environments, verify the operating system locale is same as the Process Engine locale, and verify the LC_TIME and LC_MESSAGES are set to “C” or an equivalent ISO-8859 locale before running Process Task Manager in Application Engine or Workplace XT. Failing to do so can result in character corruption and application failure.

Restriction: Application Engine or Workplace XT locales must match the Process Engine UNIX locale when starting Process Task Manager.

Limitations on installing in a non-English environment

There are certain limitations on installing IBM FileNet P8 in non-English environments.

Oracle WebLogic Server

WebLogic Server 8.1 is not supported on a Turkish operating system.

Application Engine, Workplace XT, or Process Engine Task Manager for UNIX

If Application Engine, Workplace XT, or Process Engine is installed on a UNIX platform, Process Task Manager (vwtaskman), takes too long to launch when run under a UTF-8 locale using remote CDE shells such as xWindows (exceed, xManager, etc). Verify that the xWindows application can handle Unicode fonts to

fix this problem. The command-line interface functionality is not affected.

Publishing

When the watermark is checked in the Publishing Style Template Manager on the Content Engine server, double-byte documents cannot be published. This is a known independent software vendor issue.

Process Engine Modeler

To import a user defined XSD file containing non-English characters in Process Designer, run the following command to convert characters to the Unicode encoding format to match a UNIX Process Engine in a UTF-8 locale, then import the XSD file into Process Designer.

```
Java -cp pe.jar filenet.vw.toolkit.utils.FileConverter /in filename /out outfilename
```

Appendix B. IBM FileNet P8 ports

The tables below list the port numbers used by IBM FileNet P8 components.

Content Engine ports

| Port name | Default number | Notes |
|---------------------|----------------|--|
| LDAP | 389 | This port is on the directory server, specified on the Content Engine server for authentication. |
| SSL | 636 | This port is on the directory server, specified on the Content Engine for authentication through SSL. |
| LDAP Global Catalog | 3268 | |
| WebLogic | 7001 | This port is on the WebLogic Server for Content Engine for communication with Content Engine by clients. |
| Oracle | 1521 | This port is on the Oracle database server, for communication with the database by Content Engine. |

Process Engine ports

| Port name | Default number | Notes |
|----------------------------|----------------|---|
| SMTP (E-mail Notification) | 25 | This port is on the SMTP server. |
| Oracle | 1521 | This port is on the database server, for connection between Process Engine and Oracle. |
| TMS | 32768 (TCP) | Task Manager service. TM_daemon listens for requests from initfsw running on the same or a different system analogous to COR_Listen listening for RPCs. |
| COR | 32769 (TCP) | Courier service. COR_Listen listens on this port for incoming RPC requests. |
| NCH | 32770 (UDP) | nch is the NCH daemon. NCH_daemon listens on this port. Pre-4.1.2 listened for broadcasts, etc. Post 4.1.2, only listens for old print servers to verify NCH is up. |

| Port name | Default number | Notes |
|---|----------------------------|--|
| fn_snmpd | 161 (UDP) | IBM FileNet Image Services Simple Network Management Protocol Daemon. It listens for SNMP requests from the native OS snmp daemon. The native snmp daemon listens on this port and communicates with fn_snmpd through other local port. fn_snmpd does not listen on this port. |
| fn_trapd | 35225 (UDP) | fn_trapd is the IBM FileNet Image Services trap daemon, which listens for notifications of the end of Image Services background processes running on the server and sys_logs information. |
| Rules Listener | 32774 (TCP/IP) | This port is on the Process Engine server. |
| Process Engine Communication Port (IOR port) | 32776 | This port is on the Process Engine server. |
| BPM Web Services Reliable messaging client port | 32767 (TCP) 32767 (UDP) | This port is on the Process Engine server. |
| Registry Port | 32771 | This port is on the Process Engine server and is used by the Component Manager. |
| Event Port | 32773 | This port is on the Process Engine server. |
| Process Engine Broker Port | 32777 | This port is on the Process Engine server, and is for the vworbbroker process handling RPCs from the PE API clients. |

Application Engine ports

| Port name | Default number | Notes |
|--------------|----------------|---|
| WebLogic | 7001 | This port is on the WebLogic Server for clients to connect to Application Engine. |
| WebLogic SSL | 7002 | This port is on the WebLogic Server for clients to connect to Application Engine through SSL. |

| Port name | Default number | Notes |
|---|--|--|
| Process Engine (RMI) | 32771 | <p>This port is on the Windows Process Engine server for Process Task Manager to communicate with the Windows Process Engine Services Manager.</p> <p>This port is on the Application Engine server for Process Task Manager to communicate with the Component Managers and the Windows Process Application Engine (or Workplace XT) Services Manager.</p> |
| Component Manager (Event Port) | 32773 | This port is on the Application Engine server, and is used when the Component Manager (running on the Application Engine Server) is configured to be triggered by events, instead of polling. |
| Web Services Reliable Messaging Client Port | xxxx - assigned by the client on the Application Engine Server, for a particular Component Manager instance. | This port (or a number of ports, one per Component Manager instance) is on the Application Engine server to respond to the WS-Reliable Messaging requests. |

Autonomy K2 search engine ports

| Port name | Default number | Notes |
|---------------------------------|---------------------------------|---|
| K2 Master Administration Server | 9950 (default) | This port is located on the machine that hosts the K2 Master Administration Server machine. It is the software port for K2 Master Administration Server. |
| K2 Dashboard | 9990 (default) | This port is located on the machine that hosts the K2 Master Administration Server. It is used specifically for the URL access for the K2 Dashboard. |
| K2 Index Server | 9960 - 9979 (recommended range) | This port is located on the machine where you created the K2 Index Server service. It is for communication between the K2 Index Server and P8 Content Engine. |

| Port name | Default number | Notes |
|---------------------------|---------------------------------|--|
| K2 Broker Server | 9900 - 9909 (recommended range) | This port is located on the machine where you created the K2 Broker Server service. It is for communication between the K2 Broker Server and P8 Content Engine. |
| K2 Server (Search Server) | 9920 - 9949 (recommended range) | This port is located on the machine where you created the K2 Server (Search Server) service. It is for communication between the K2 Server (Search Server) and Content Engine. |
| K2 Ticket Server | 9910 - 9919 (recommended range) | This port is located on the machine where you created the K2 Ticker Server service. It is for communication between the K2 Ticket Server and Content Engine. |

Rendition Engine and Content Engine ports for Liquent

| Port name | Default number | Notes |
|--------------------|----------------------------|--|
| Liquent input port | 2867 (COM Repository only) | This port allows for distributed processing of jobs on the Rendition Engine servers when there are more than one Rendition Engine server. A Rendition Engine server or Rendition Engine client (Content Engine Publishing server) on which a job is submitted makes the load balancing decision on where to run the job. If the job is sent to another Rendition Engine server for execution then it uses the other Rendition Engine server's input port to do so. |

| Port name | Default number | Notes |
|----------------------------|----------------|--|
| Liquent notify port | 2868 | This port is also related to the distributed processing of jobs on the Rendition Engine servers. When a job is sent to another Rendition Engine server then the originating Rendition Engine server or Rendition Engine client (Content Engine Publishing server) is notified on this port by the other Rendition Engine Rendition Engine server when the other Rendition Engine Rendition Engine server has finished processing the job. One case of this is when the Rendition Engine Rendition Engine server is notifying the Rendition Engine client (Content Engine Publishing server) that a conversion job has completed. |
| Liquent event port | 2869 | This port is used by each Rendition Engine server to send or receive events. These events are used by the Rendition Engine servers to keep each other informed of current activities. The Rendition Engine or Liquent Domain Manager uses these events for the job status display. |
| Liquent admin port | 2870 | This port is used by each Rendition Engine server for internal administrative functions. It's primary use is for each Rendition Engineserver to publish its current activity statistics (for example, how busy the business services are) for use by other Rendition Engine servers to make load balancing decisions. |
| Liquent file transfer port | 2871 | This port is used for transferring source, temporary, and result files between Rendition Engine servers as well as between Rendition Engine servers and Rendition Engine clients (Content Engine Publishing servers). |

| Port name | Default number | Notes |
|------------------------|----------------|--|
| Liquent job queue port | 2872 | This port is for job queuing by the render business service on all Rendition Engine servers. |
| Oracle | 1521 | When the selected Rendition Engine database is Oracle, then this port is used by the Rendition Engine servers and the Rendition Engine clients (Content Engine Publishing server) to access the Rendition Engine database. |

Tip: If the Liquent port number assigned to the Rendition Engine/Liquent software conflicts with the port number required by another application or service that runs on the Rendition Engine server or the Content Engine Publishing server, then the default values can be changed in the Rendition Engine/Liquent Domain Manager. The above port numbers are the default values set by the Rendition Engine installer for the COM Repository in the Rendition Engine/Liquent domain, which represents the Rendition Engine server itself. When a Content Engine Publishing server is configured to point to the Rendition Engine server, one must create a "Java Repository" in the Rendition Engine/Liquent domain that represents the Content Engine Publishing server as a Java client to the Rendition Engine server, and this repository will also have the same default port number values.

The database port number is not specified directly by the Rendition Engine server's Liquent software, but it is specified in the Rendition Engine client/Content Engine Publishing server's Rendition Engine Connection (for the JDBC connection from the Content Engine Publishing code to the Rendition Engine database server).

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating

platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml> (www.ibm.com/legal/copytrade.shtml).

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, and service names may be trademarks or service marks of others.

Index

A

- accounts
 - Application Engine 37
 - Content Engine 34
 - Content Engine application server 58
 - Content Search Engine 38
 - Process Engine 36
 - specifying 33
- Application Engine
 - configuring Solaris 23
- ASA
 - installation tasks 57
- authentication
 - Sun Java System Directory Server 31

C

- configurations
 - Baseline with options 2
 - demo 3
 - developer 3
 - sample 1
- Content Engine
 - configuring remote access protocol 28
 - remote file access protocols 26
- Content Search Engine
 - preparing storage areas to support 28

D

- database
 - configuring automatic transaction processing (Oracle) 49
 - creating 46
 - dedicated (Oracle) 42
 - environment variables (Oracle) 49
 - GCD (Oracle) 47
 - installing (Oracle) 42
 - object store (Oracle) 47
 - Process Engine tablespaces (Oracle) 48
 - shared (Oracle) 42
 - storage area 24
 - verifying connection 56
 - XA transactions (Oracle) 50
- database accounts
 - creating 41
- database connection
 - verifying connection for Oracle 56
- DB2
 - JDBC drivers 58
 - WebLogic 58
- DBA
 - installation tasks 39
- Decru 24
- default time mask
 - Solaris Process Engine 20

- deployment
 - planning 3
 - planning application server 4
 - planning Content Engine 3
 - planning Content Engine multi-server on WebSphere and WebLogic 5
- directory servers
 - Sun Java System Directory Server 31
- documentation server
 - installing 61
- domain 60

E

- encrypted NTFS devices 24
- Encryption products for storage 24
- extensible authentication ports 71

F

- file storage area
 - configuring 25
 - defined 24
- fixed storage area 24

H

- hostname 18
- hosts file 18

I

- IBM Encryption expert 24
- installation
 - ASA tasks 57
 - DBA tasks 39
 - ITA tasks 10
 - planning 1
 - planning and preparing 1
 - SA tasks 29
- installation and upgrade worksheet 7
- iPlanet, see Sun Java System Directory Server 31
- IPv4 18
- IPv6 18
- ITA
 - installation tasks 10

J

- JDBC drivers
 - install for DB2 58
 - install for SQL Server 58
 - WebLogic 58
 - install for Oracle 58

K

- Kerberos 29

L

- LANG settings
 - Solaris Process Engine 20

N

- network name 18

O

- Object Request Broker(ORB)
 - application server-based PE clients 61
- operating system accounts
 - creating 13
- Oracle
 - character sets 46
 - client install 55
 - dedicated database 42
 - install for Content Engine 42
 - install for FileNet P8 platform 43
 - install for Process Engine 42
 - installing for FileNet P8 components 45
 - JDBC drivers
 - WebLogic 58
 - shared database 42
- Oracle client
 - installing 55

P

- password complexity (Oracle) 50
- ports 71
- Process Engine
 - configuring Solaris 19
 - configuring UNIX 17
 - configuring UNIX servers 17
 - database accounts, creating 41
 - enable port 32777 19
 - operating system accounts, creating 13
- Process Engine (Solaris 10)
 - kernel limits 21
- Process Engine (Solaris 9)
 - kernel limits 20

R

- realm 29
- roles
 - definition of installation 5

S

SA

- installation tasks 29

security

- install considerations 29
- specify accounts 33

Single Sign-On 29

SQL scripts (Oracle)

- changes made by 53, 54
- changing default passwords 54
- overview 52
- running 50
- running manually 51, 54

SQL Server

- JDBC drivers 58
- WebLogic 58

SSO

- Kerberos 29
- planning considerations 29
- requirements 29

Storage areas 24

Sun Java System Directory Server 31

- setting resource limits for directory server 32
- setting resource limits for users 32

Sun ONE, see Sun Java System Directory Server 31

T

time mask

- Solaris Process Engine 20

U

UNIX

- configuring for FileNet P8 components 16

V

Vormetrics 24

W

WebLogic

- configure for Content Engine 58

WebLogic Server

- configuring for Application Engine 60
- configuring for Workplace XT 60

Workplace

- avoid logon failures in multi-forest domains 58

worksheet

- autofiltering and sorting 9
- running the customize macro 8
- using the installation and upgrade worksheet 7



Program Number: 5724-R76
5724-R81

GC19-2826-01

